

[54] **SERIAL LAYERED MEDICAL NETWORK**

[75] **Inventor:** Michael Fischer, San Antonio, Tex.

[73] **Assignee:** Nellcor Incorporated, Pleasanton, Calif.

[21] **Appl. No.:** 829,146

[22] **Filed:** Jan. 31, 1992

[51] **Int. Cl.** H04J 3/24; H04L 29/02;
H04L 12/56

[52] **U.S. Cl.** 370/94.1

[58] **Field of Search** 370/94.1, 60, 85.15,
370/58.1, 56, 112, 92, 85.13, 42, 83, 85.6

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,375,097	2/1983	Ulug	370/94.1
4,423,414	12/1983	Bryant et al. .	
4,495,493	1/1985	Segarra et al. .	
4,549,297	10/1985	Nishimoto .	
4,574,284	3/1986	Feldman et al. .	
4,603,416	7/1986	Servel et al.	370/94.1
4,680,581	7/1987	Kozlik et al. .	
4,692,918	9/1987	Elliott et al. .	
4,706,080	11/1987	Sincoskie .	
4,930,122	5/1990	Takahashi et al.	370/94.1
4,941,089	7/1990	Fischer .	
4,947,390	8/1990	Sheehy	370/85.13
5,007,043	4/1991	van den Dool et al.	370/94.1
5,010,546	4/1991	Kato	370/94.1
5,050,166	9/1991	Cantoni et al.	370/61
5,067,123	11/1991	Hyodo et al.	370/58.1
5,113,392	5/1992	Takiyasu et al.	370/85.15
5,140,584	8/1992	Suzuki	370/60
5,173,897	12/1992	Schrodi et al.	370/60
5,214,642	5/1993	Kunimoto et al.	370/82

OTHER PUBLICATIONS

"Internetworking and Addressing", Gene White, Chapt. 5, pp. 101-125, McGraw-Hill, Inc., 1992.

"Architecture of a Comprehensive Radiologic Imaging Network", H. K. Huang et al., IEEE Journal on Selected Areas in Communications, Sep. 1992, vol. 10, No. 7, ISACEM (ISSN 0733-8716).

Primary Examiner—Wellington Chin

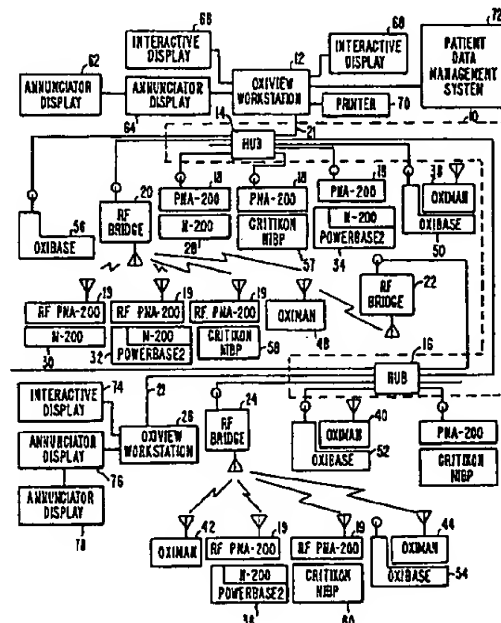
Assistant Examiner—Ajit Patel

Attorney, Agent, or Firm—Townsend and Townsend and Crew

[57] **ABSTRACT**

A network or telemetry system which allows virtual services at the application or presentation layer to communicate with other virtual services without regard to the physical interconnections. Each message, called a parcel, includes the information to be transmitted along with a virtual address header. The parcel is provided to a gateway, which inserts the parcel without modification into a packet with address information for the physical through session layers in the packet header. The packet is then transmitted to another network node, which receives and delivers the unmodified parcel to the addressed destination virtual service. A number of parcels from the same or different virtual services can be packed into a single packet for transmission from the gateway in cases where these parcels are all directed to virtual services at the same destination node. Once a session is established, such as between a gateway and a workstation, virtual services at the gateway node and the workstation can communicate with each other without requiring a lot of header overhead for each transmission. Instead, the session need simply be identified. Each gateway typically has one session at a time, but a workstation can support up to 64 sessions simultaneously.

12 Claims, 8 Drawing Sheets



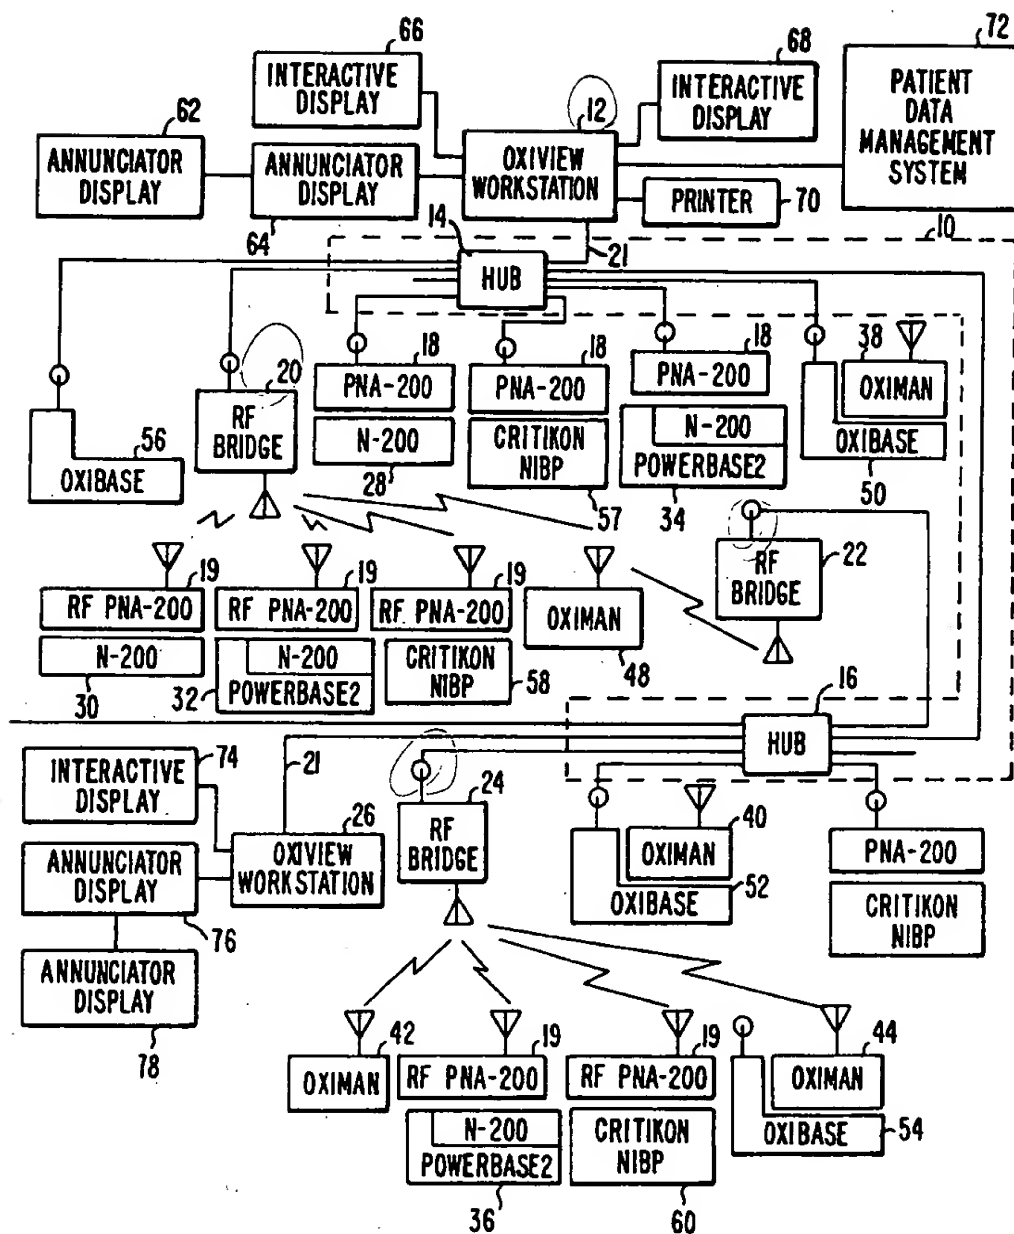


FIG. 1.




FIG. 2.




FIG. 3.




FIG. 4.




FIG. 5.




FIG. 6.




FIG. 7.

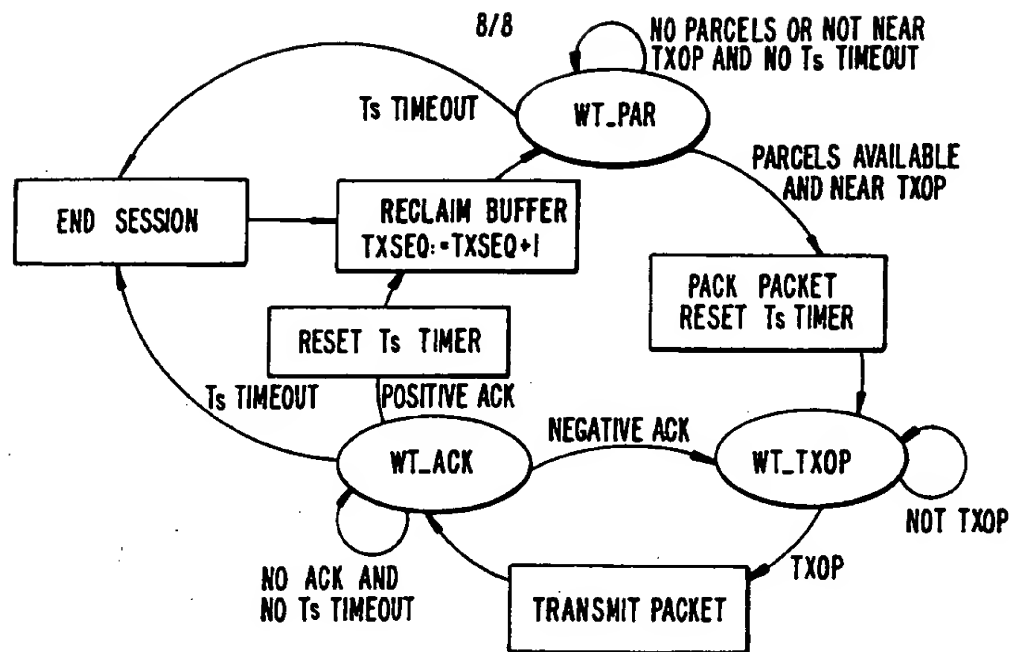


FIG. 8.

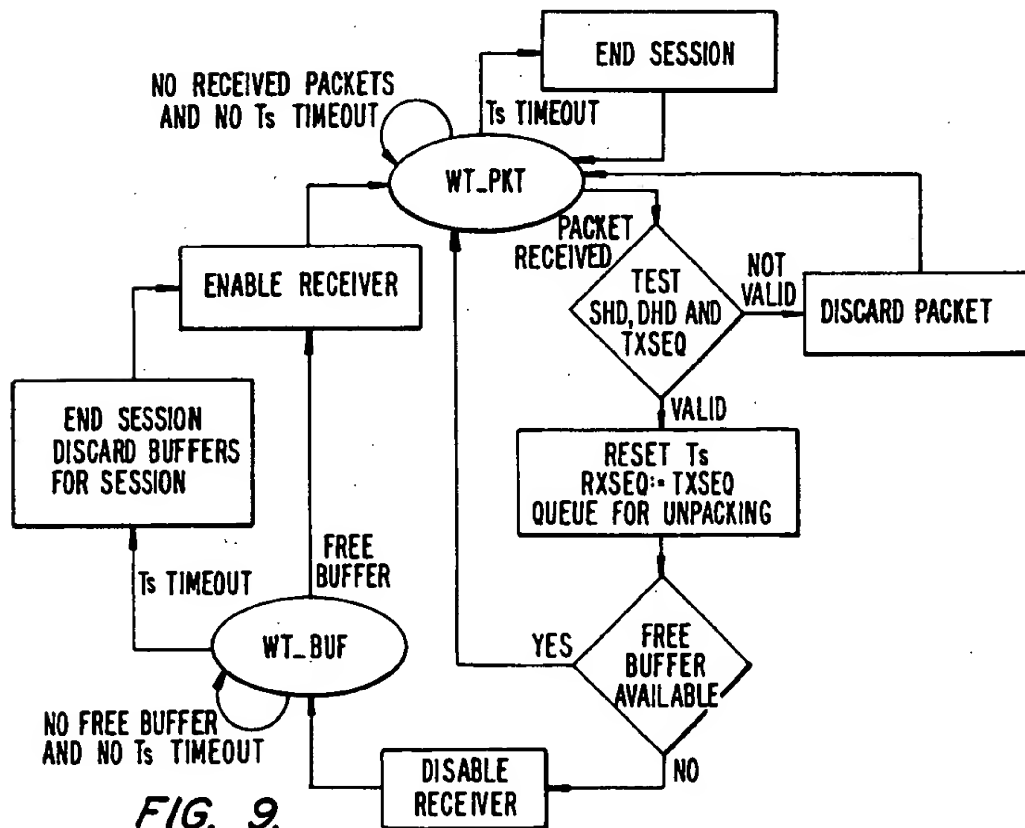


FIG. 9.

SERIAL LAYERED MEDICAL NETWORK

BACKGROUND OF THE INVENTION

The present invention relates to serial data communication networks, and in particular to networks for interconnecting medical instrumentation.

In a typical computer network, computers are connected together over a communication medium. Each computer has its own, unique physical address which is used for identifying both the source and the destination of any transmission. Data and other information is typically sent in packets, with each packet containing a data field and a header setting forth the source and destination addresses, as well as other information. Different protocols exist for the header and for determining when a particular source can transmit.

In a number of fields, such as the medical field, it is desirable to be able to connect remote instruments to a central computer workstation. Typically, the instruments will gather data and have minimal processing power. The large bulk of data is typically transmitted from the instruments to the computer. In addition to the data, there may be alarm signals which need to be transmitted and immediately received.

It would be desirable to have a system optimized for network communication from a number of medical or other instruments to a central computer.

SUMMARY OF THE INVENTION

The present invention provides a network or telemetry system which allows virtual services at the application or presentation layer to communicate with other virtual services without regard to the physical interconnections. Each message, called a parcel, includes the information to be transmitted along with a virtual address header. The parcel is provided to a gateway, which inserts the parcel without modification into a packet with address information for the physical through session layers in the packet header. The packet is then transmitted to another network node, which receives and delivers the unmodified parcel to the addressed destination virtual service.

A number of parcels from the same or different virtual services can be packed into a single packet for transmission from the gateway in cases where these parcels are all directed to virtual services at the same destination node. Once a session is established, such as between a gateway and a workstation, virtual services at the gateway node and the workstation can communicate with each other without requiring a lot of header overhead for each transmission. Instead, the session need simply be identified. Each gateway typically has one session at a time, but a workstation can support up to 64 sessions simultaneously.

For example, a gateway with a pulse oximeter attached may establish a session with a workstation. The pulse oximeter would provide virtual services for real time data streams for oxygen saturation values, ECG values and pulse values. A separate virtual service called trend service would periodically store real time data for subsequent retrieval. These would communicate with virtual services in the workstation over a single established session. The oxygen saturation and ECG may communicate with a display control virtual service at the workstation, while the pulse value service communicates with an annunciator service at the workstation, for instance. The workstation can simultaneously carry on other sessions with other pulse oximeters

or other instruments. A single session may last the duration of a patient's stay in a hospital room.

Unlike the prior art, where separate sessions would typically be needed for transmissions between each pair of virtual services, the present invention supports transmissions between multiple virtual services in a single session. This eliminates the need for each instrument or virtual service to have a large amount of computing power to support its own session. By sharing a session, less overhead in the form of computing power to support communication to and from multiple virtual services is required; while still permitting the virtual services to be unaware of data handling during the communication process.

The parcels can be of varying size and number. Each parcel includes precedence information in the header which indicates the relative delivery importance of the information contained in the parcel. For example, an alarm indication parcel would have a highest precedence level, while real-time data would have lower precedence (with further distinctions between types of data: general data would be higher precedence than detailed data which would be higher precedence than stored data that could be resent if necessary). The gateway transmits the highest precedence level parcels first. If buffer space at the gateway is exhausted, parcels having the lowest precedence level are discarded. If only high precedence parcels are present, older parcels are overwritten with the newer parcels from the same source.

Each virtual service sends parcels to the gateway with precedence information in the parcel header. The parcel header also identifies the length of the information field. The information field can contain data, a command requesting an action, or a reply to a request. If the information field contains data, a sequence number is included indicating the order in which the parcel was generated.

Each gateway has a table for indicating the location of virtual services local to that node and the internal addressing required to deliver parcels to those virtual services. This parcel routing is done transparent to the virtual service itself. The gateway also has a buffer for temporarily storing parcels while they are waiting to be multiplexed into a packet. The packet header identifies the number of parcels included and the overall length of the packet information field containing the parcels. The packet header also contains source and destination handles identifying the physical source node and destination node, as well as the particular session (which is useful for nodes that support a plurality of sessions). A sequence number is included to identify the packet for detection of packets delivered more than once by the physical network hardware.

The present invention detects missing data at the application layer for each service. This is done with the assumption that there is a reliable physical/MAC layer underneath. In this context, "reliable" means that the physical and MAC layers are (when communication is possible) incapable of indicating packet delivery without the packet having successfully reached the destination node (at the MAC layer). In order to do this, the MAC layer, by necessity, is capable of delivering the same packet twice.

For fuller understanding of the nature and advantages of the invention, reference should be made to the ensuing detailed description taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a system using a protocol according to the present invention;

3

FIG. 2 is a block diagram of a peripheral network adaptor (PNA) of FIG. 1;

FIG. 3 is a block diagram of a hub of FIG. 1;

FIG. 4 is a block diagram of a workstation of FIG. 1;

FIG. 5 is a block diagram of the intelligent network adapter (INA) of FIG. 4;

FIG. 6 is a diagram of the ISO layers used by the present invention;

FIG. 7 is a diagram showing the fields of the parcels and packets of the present invention;

FIG. 8 is a diagram of a transmitter state machine; and

FIG. 9 is a diagram of a receiver state machine.

DESCRIPTION OF THE PREFERRED EMBODIMENT

FIG. 1 is a block diagram of one embodiment of a system using the protocol of the present invention. An Oxinet2 network 10 connects workstations 12 and 26 to a large number of medical instruments. The network is physically configured, for wiring convenience and other reasons, as a number of lines which are physically interconnected through hubs 14 and 16. Communication over network 10 is controlled by gateways connected to each node. Each of workstations 12 and 26 includes an internal intelligent network adapter (INA) which acts as a gateway. Each of the other instruments connected to the network either contains its own internal gateway or is connected through a gateway called a peripheral network adaptor (PNA) 18. A number of instruments are connected to the network through a radio frequency link that operates between radio-equipped gateways 19 and radio frequency (RF) bridges 20, 22 and 24 that serve as radio hubs and connect the RF network to the wired network.

Among the instruments shown in FIG. 1 are pulse oximeters 28 and 30. Other pulse oximeters 32, 34 and 36 include a power base with wave form display and printer capabilities. Personal monitors 38, 40, 42 and 44 monitor heart rate, respiration, oxygen saturation and ECG waveforms. Monitors 38, 40 and 44 include a base unit 50, 52, 54, respectively, which provides a bedside waveform display capability. A display 56 is shown separately hooked up to hub 14. Also shown are blood pressure monitors 57, 58 and 60. The PNA gateway devices 18 and RF PNA gateway devices 19 allow older, existing instruments out in the field to be adapted for communication over the network of the present invention.

Each workstation has other elements coupled to it. Workstation 12 has a pair of annunciator displays 62 and 64, a pair of interactive displays 66 and 68, a printer 70, and is coupled to a hospital's patient data management system 72. Workstation 26 is connected to a single interactive display 74 and a pair of annunciator displays 76 and 78. Annunciator displays only present alert conditions. The interactive displays permit a user to request the display of alerts, status, data, trends and waveforms.

FIG. 2 is a block diagram of a peripheral network adaptor or gateway 18 or 19 of FIG. 1. The connection to its associated instrument is through a serial port 210. The connection to the network is through a network interface 212 for wired gateways 18 or a radio interface 214 for RF gateways 19.

The PNA has a microcontroller 218 which is connected to push buttons 220, status LEDs 222, and EEPROM memory 224 that is used to store configuration parameters. A separate

4

flash EPROM memory 226 is used to hold operating firmware. An SRAM memory 228 acts as the buffer for the PNA, storing parcels which are assembled into packets and holding trend data for up to 24 hours for readout by the workstation upon request. The PNA also includes a clock/calendar circuit 230, and a power supply 232 with a connection to an external AC adapter 236. A battery 234 is used to maintain the contents of the clock/calendar and SRAM circuits for at least 24 hours without AC power.

For a wired connection to the network, a separate ARCNET controller circuit 238 handles the network transmissions, implementing the ARCNET MAC layer. "ARCNET" is a registered trademark of Datapoint Corp. ARCNET is a widely used LAN described in the ARCNET Designers Handbook, Datapoint Corporation, 2nd Ed., 1988, order no. 61610. ARCNET is also covered by a proposed ANSI Standard, draft rev. 1.6, 1-5-91, available from the ARCNET Trade Association, 3365 North Arlington Heights Road, Suite J, Arlington Heights, Ill. 60004.

Microcontroller 218 controls the assembling of parcels into packets for transmission. If a radio interface is used, RF network interface 240, radio modem 42, and RF antenna 244 transmit to a RF bridge as shown in FIG. 1. The RF bridge contains a microcontroller similar to microcontroller 218 and an ARCNET controller similar to controller 238. Also shown in the radio interface 214 is an infrared location interface 246. This is used to receive signals from an infrared transmitter in a hospital room to indicate the location of the instrument.

PNA microcontroller 218 acquires data from an instrument over EIA-232 link 210. The link is configured to be in a mode compatible with the link of the instrument. A peripheral transaction server (PTS) in microcontroller 218 is programmed to receive data from the instrument. The data acquired from the instrument is stored in a circular buffer in the microcontroller where the raw instrument data is assembled and routed to a buffer position in SRAM 228. EEPROM 224 is a configuration memory which holds session and initialization parameters. The PNA microcontroller executes firmware stored in flash EPROM 226 that includes functions that translate instrument-specific data formats into the uniform parcel formats used by the workstations. This firmware also maintains a record of actual sensor data each 10 seconds in a special area of SRAM 228 allocated for such trend storage, and permits these trends, extending back up to 24 hours, to be read out by the workstation.

FIG. 3 is a block diagram of hub 14 of FIG. 1. The network connections are made through the transceivers for ports 1-16 as shown on the right side of the figure. These ports connect to the workstations and the various instruments and PNAs as shown in FIG. 1. A logic block 310 interconnects the lines depending upon who is transmitting to whom. The logic is controlled by an ARCNET controller 312 and a separate microcontroller 314. There is also provided a signal retiming (delay line) logic 316 and network address detection logic 318. Status LEDs 320 are also provided, along with a power supply 322. The hub listens for transmissions from the different ports. When a transmission is detected, the hub will receive on that port and re-transmit on the remaining ports.

FIG. 4 is a block diagram of a workstation 12 of FIG. 1. The connection to the ARCNET network 10 is through a intelligent network adapter (INA) 410. The workstation is operated under the control of a main, root processor 412 with main memory 414. An internal bus 416 interconnects

the elements under the control of a system bus controller 418. A graphics adapter 420 provides data to a video display or a touch services board 422. An SCSI controller 424 provides an interface to disk memory 426. Different virtual services in workstation 12 communicate with virtual counterparts in the instruments attached to the network using the protocol of this invention. An example of a virtual service is a program running on the workstation that performs a particular function.

FIG. 5 is a block diagram of the INA 410 of FIG. 4. The connection to the network is through an ARCNET controller 510 and transceiver 512. The connection to the internal bus 416 of the workstation is through a memory space interface and relocation register 514 and an I/O space interface 516. The INA operates under the control of its own processor 518, and includes EPROM memory 520, SRAM memory 522, and EEPROM memory 524. Also included is a watchdog timer 526, switches and LEDs 528 and an optional serial interface 530 for debug purposes. INA 410 acts as the gateway for workstation 12, assembling parcels into packets for transmission, and deassembling received packets.

Transmissions primarily take place from the instruments to the workstation, and from the workstation to the instruments. An instrument sending data or other information transmits a parcel to its associated gateway. The gateway receives several parcels from various instruments or the same instrument and packs them into a packet in the order of arrival. If more parcels arrive between packet transmission opportunities than will fit into a packet, precedence is used to cause the most important parcels to be packed into packets first. The packet is then transmitted to the remote workstation as discussed in detail below. The instrument need only provide the virtual service address of the application layer program in the workstation, with the intermediate levels through physical layer transmission being taken care of by the gateway and the network.

The gateway will examine the precedence information in the parcels it receives and send alarms from its associated instruments first, in accordance with the precedence rules discussed below. The parcels awaiting transmission in the gateway's buffer are overwritten in accordance with the precedence rules in the event of buffer exhaustion. This prevents a real time data waveform from preventing an alarm from getting through, for instance.

1. Definitions

The meanings of some terms used are defined below. Terms being defined are presented in bold, while the first references to terms whose definitions appear subsequent to their initial usage are presented in *italics*.

Oxinet2 is the name of a network using the protocol of the present invention.

The **Oxinet2**

protocol is a definition of communication over Oxinet2 networks.

An **Oxinet2**

network is a set of entities communicating among each other, via a serial interconnection medium, using packets conforming to the Oxinet2 network protocol.

The elements that use the Oxinet2 network communication facilities are

services, which are application layer entities that exchange parcels with other services.

Each Oxinet2 network is comprised of one or more

segments, which are subsets of the network that communicate via a single instance of a particular link layer facility. Examples of link layer facilities currently used for Oxinet2 segments are

ARCNET (a registered trademark of Datapoint Corporation), used to provide 2.5 Mbps of raw transfer bandwidth over coaxial cable to locations where wired interconnection is available, and the

Oxinet2 RF network, used to provide a spread-spectrum radio link to mobile devices and/or locations where wired interconnection is not available.

Each segment using the Oxinet2 RF network operates on a particular channel, which is a subset of the available portion of the RF spectrum, using a particular frequency range and spreading sequence, to offer 200 Kbps of raw data transfer bandwidth.

Communication between Oxinet2 network segments takes place through

bridges, which are clusters that forward packets between network segments, performing any necessary buffering and conversion of frame formats. Bridges direct packets according to

mutes, that associate network-specific addresses on each network segment attached to the bridge with the unique station ID in the packet header.

One type of bridge used for Oxinet2 networks is an RF bridge, which forwards packets between an ARCNET segment and an RF channel that constitutes an Oxinet2 RF network segment.

The session administrator (SA) is a service facility, typically implemented in software on an Oxiview workstation, that

receives requests to initiate communication activities; determines whether to create sessions in response to these requests,

determines whether each new session is a reconnection of a preexisting session, and

provides the necessary information to permit bridges to determine the appropriate route for the communication activity.

Communication between Oxinet2 networks takes place through

gateways, which are stations that forward parcels between the network and virtual services, performing any necessary multiplexing and demultiplexing functions.

The basic addressable units on a Oxinet2 network are stations, each of which are uniquely-identified physical entities from among all manufactured equipment which may be attached to, and exchange information over, an Oxinet2 network.

One or more physically connected stations constitute a cluster, which is a component, or a set of interconnected components, attached to an Oxinet2 network. Examples of clusters include

Oxiview workstations (including their attached peripheral devices), each of which is a single station;

PNA-200 peripheral network adapters (including their attached peripheral devices).

The term

local refers to facilities or services which are internal to the entity of reference (typically a cluster), whereas

remote refers to facilities, services, or entities which are external to the entity of reference.

Satellite refers to PNA-200s in contrast to Oxiview workstations and RF Bridges. For example a

satellite RF station, is a station with a transceiver for the Oxinet2 RF network. The satellite RF station is remote from the reference points of Oxiview workstations and RF bridges.

The implementation of network facilities is described in

streams are totally arbitrary as viewed from the presentation layer.

A goal of Oxinet2 network communication is to minimize the

observable delay, which is the time between a user's perception of the occurrence of an event (such as an alarm condition) at a satellite monitor and the reporting of the occurrence of that same event to a care provider at the connected workstation.

Each Oxinet2 service is identified by a

service name, that is assigned globally.

Communication between services is directed using a three-component address that consists of a

destination tag (DTAG),

recipient service ID (RID), and

action code (ACT).

The 8-bit tag identifies an instrument, workstation, or bridge. Both a

source tag (STAG) and a

destination tag (DTAG) are included in each parcel.

Tag values are assigned globally.

The 8-bit service ID is a virtual label that uniquely identifies a particular service within the context of the module identified by the associated tag. Both an originator service ID (OID) and a recipient service ID (RID) are included in each parcel, with the exception of a few globally-assigned service IDs, the service ID values are assigned dynamically by the session manager whenever a satellite cluster is initialized.

The 16-bit action code (ACT) uniquely identifies the action to be performed by the recipient of the parcel, and is used, within the recipient module to direct the parcel to the task appropriate for the designated action.

A partner is identified by a

32-bit handle, which is a globally unique identifier value.

The handle is subdivided into a

24-bit station ID (STID), which uniquely identifies the station at which the partner exists; and an

8-bit session number (SSN), assigned uniquely by the Oxinet2 network layer at the designated station.

In the case of multi-station clusters, the station ID in the handles used for session communication to and from that cluster identifies the station serving as the session manager for the cluster.

In terms of facilities provided by the Oxinet2 protocol, a workstation is a single station as far as any satellite station is concerned.

The station ID consists of an

8-bit station type (STP), which identifies the kind of station and a

16-bit serial number (SRL), unique within each station type value. More than one station type value may be assigned for a single kind of station if the range of serial numbers for that kind of station is expected to exceed 64K. Zero is not a permissible value for a serial number.

In the case of multi-station, modular instruments, multiple station IDs exist, one for each internal module.

The station ID in the handles used for session communication to and from the instrument identifies the gateway module of that instrument.

An 8-bit instrument type (ITP) is maintained on the module that provides user interface (UIF) service for

the instrument. In order to allow instrument type values to directly be used as instrument tag values, instrument types are restricted to values between 1 and 223, inclusive.

The instrument ID (IID) is the 32-bit value obtained by concatenating the instrument type with the station ID of the module that provides user interface service for the instrument. The instrument ID may be used to uniquely identify an instrument for all short-duration activities, including the determination by a workstation that the same instrument has reconnected via a different gateway. However, the instrument ID may be altered during repair of the instrument, so long-duration activities, such as physical asset tracking, should be done using the serial number recorded on the instrument enclosure, not the instrument ID.

With respect to any parcel or packet transferred between stations, the

source is the station at which the parcel is transferred from the presentation layer, or the packet is transferred from the transport layer, to lower-layer facilities. While the destination is the station at which the parcel is transferred to the presentation layer, or the packet is transferred to the transport layer, from lower-layer facilities.

There is only one source and one destination for each parcel or packet, independent of the number of intermediate stations, such as bridges or gateways, through which the parcel or packet passes enroute.

In the presence of bridges and/or gateways, there is an important distinction between the source and destination, which are distinguished by operations that occur at, or above, the transport layer; and the transfer facilities of the underlying MAC layer. Significant MAC layer operations are performed by the

initiator, which is the station that transmits a data packet frame onto a network segment and the

target, which is the station that receives, and may acknowledge, the data packet frame sent by the initiator.

In an inter-segment transfer through a single bridge:

the source is also the initiator of the data packet frame on the first network segment, but the target of that data packet frame is the bridge, which is not the destination; and

the bridge is the initiator, but not the source, of the data packet from on the second network segment, where the target of that data packet frame is also the destination.

At the physical layer

the terms transmit and receive apply to the stations that physically transfer any type of frame to and from the network medium. Therefore:

An initiator transmits a data packet frame to a target, which receives the data packet frame.

After receiving the data packet frame, the target may transmit an acknowledgement frame to the initiator.

In descriptions of programmatic interfaces the terms

upcall refers to a procedure call made by a handler for a lower layer of the Oxinet2 protocol, such as a MAC-layer packet driver, to a higher-layer facility, such as a transport-layer transport control state machine, to inform that higher-layer facility of the occurrence of an event (typically the arrival of a packet or parcel); and

downcall refers to a procedure call made by a facility for a higher layer of the Oxinet2 protocol, such as a transport-layer transport control state machine, to a

of this system code in all Oxinet2 packets on ARCNET permits the Oxinet2 protocol to coexist with other, simultaneous usage of the same physical and MAC layers by other protocols serving other software and users. However, despite this use of an ARCNET system code, it is not intended that ARCNET segments of an Oxinet2 network be shared for any purpose other than the interconnection of Oxinet2 devices.

Network Layer

The network layer 618 is concerned with the establishment and maintenance of connections between logical entities on the network. These functions involve:

1. Routing, in which the session administrator service at the workstation establishing the session configures the bridge(s) that connect the appropriate network segments to forward packets appropriately;
2. Roaming, in which the session administrator service at the workstation handling the session detects that communication with a satellite RF station that was lost via one bridge has been reestablished via a different bridge, and reestablishes both the route and the transport layer communication appropriately; and
3. Packet validation, in which addresses are recognized and filtered for packets being handled by or for the transport layer.

The session administrator service is implemented by software on the Intelligent Network Adapters (INAs) installed in workstations. The packet validation functions are implemented on all stations, either by firmware or by driver-level software.

Transport Layer

The transport layer 620 is responsible for moving the data stream between logical entities at various points on the network. Packetization, retransmission due to MAC layer errors, and routing control are transparent to users of the transport layer. The transport layer is provided by transport control state machines implemented in software on gateway, (hub) processors, INAs, and RF bridges.

Session Layer

Transport services take place in the context of sessions, which are virtual circuits for the exchange of packets between particular pairs of communicating workstations and gateways. Sessions are established using the network layer services, and operate through direct access from the transport control state machines to the MAC layer packet drivers once the session is established. Session connection and disconnection functions are provided as preprogrammed functions on the gateway processors in satellite interfaces to network segments. Connect request handling is implemented in software at workstation interfaces to network segments. RF bridges operate at the network and MAC layers, and do not participate in sessions.

The session layer also handles the multiplexing of parcels sent by different services into packets for transmission to the session partner, and the demultiplexing of parcels sent by the session partner for delivery to the designated destination services. Preferential handling of certain parcels, and selective discarding of parcels during periods of buffer exhaustion, is performed at this layer based on parcel precedence. The software interface to the session layer is known as a parcel driver.

The protocol definition of the present invention ends at the session layer. For completeness, the higher layers are discussed below.

Presentation Layer

Because data transfers utilize independently-generated parcels, the principal activity at the presentation layer is to

serve as an Application Program Interface (API) to the programming environment(s) used to implement the application layer functions. This API transfers parcels to and from the application layer, and can therefore provide uniformity for access to local services (within a cluster of instruments attached to a hub) and to remote services, that require use of at least one network segment, and a session, to communicate with stations external to the cluster.

Application Layer

The entities which actually perform the useful work of a system exists at the application layer. With the exception of some utility programs used to configure and maintain the network, these entities are outside the scope of this protocol.

3. Addressing

FIG. 6 shows the different fields of a parcel 710 and a packet 712. Application layer entities are identified by service identifiers (OID and RID in parcel 710), each of which are virtual labels that identify a service. Parcels 710 are sent between such services without any direct knowledge by either the originator or the recipient of the parcel regarding where the other service provider resides. This form of virtual addressing permits both the processor on which a service is performed, and the network path used to communicate with that processor, to change dynamically, while communication is in progress, without interfering with application layer communication. The operation of the protocol shields the application layer entities from any awareness of these changes. The key to implementing such transparency is the addressing structure.

Addresses are layered, in keeping with the protocol layering:

1. Service identifiers (OID, RID) and tags (STAG) DTAG) constitute presentation layer addresses in parcel 710,
2. session numbers (SSN) constitute session layer addresses in packet 712,
3. station IDs (STID) constitute network and transport layer addresses in packet 712, and
4. MAC source (SID) and destination (DID) IDs constitute MAC layer addresses.

Network Addresses

Network addresses, referred to as handles, contain the unique identifier of a session layer entity. This entity is a station in the case of all device types except workstations, or is the handler for a particular data stream within an Oxiview workstation.

A handle is a 4-byte value comprised of two fields:

1. The first 3 bytes of the handle are the station ID (STID), that provides the network layer address to uniquely identify the network interface used by the session layer entity represented by this handle; while
2. The fourth byte of the handle is the session number (SSN), that provides the transport layer address to uniquely identify the session to which this packet belongs, within the context of the designated station. Session number zero is used strictly for packets concerned with establishment of a session or with gathering network statistics from the station.

BYTE	LAYER	NAME	USED TO IDENTIFY
0	Network	STP	Station type
1-2	Network	SRL	Serial number within type
3	Transport	SSN	Session number

Action Codes

Action codes uniquely identify functions to be performed at recipient services. Action code values are used, within most modules, to direct incoming parcels to the proper internal task for processing.

Action codes are assigned globally. Within each range the specific code assignments are arbitrary, so long as they are uniquely decodable. For network control actions, the action codes and associated parcel formats are defined below under Network Operation. For instrument parcels, the recommended practice for action code assignment is as follows. The high-order byte of the action code is the service name value. The high-order two bits of the low-order byte of the action code are encoded in the same manner as the Kind bits (bits 7-6) of the parcel flags byte. The low-order six bits of the low-order byte of the action code are assigned in ascending numerical sequence, with the restriction that the same value should have the same generic meaning across different services. For example, if the value 3 is used for periodic trend update, this will apply for SpO service, respiration service, ECG service, etc.

Frame Formats

The Oxinet2 protocols are designed for usage over a plurality of physical and link layers, including ARCNET and the Oxinet2 RF Network. This section depicts the layer of the fields within Oxinet2 packets, the layout of the fields within Oxinet2 parcels, the packing of parcels within packets, the encapsulation of packets within ARCNET and RF Network MAC frames, and the encapsulation of parcels within ARCNET MAC frames.

The Oxinet2 Protocol Formats diagram in FIG. 7 depicts the layering of all Oxinet2 facilities from the link through the presentation layers.

Packet Format

All Oxinet2 packets consist of not more than 248 bytes, including a 12-byte Oxinet2 (network/transport/session layer) header, and a 10- to 236-byte information field.

Limiting the packet size to 248 bytes permits use of short packet mode on ARCNET, and keeps the total strength of RF packets, including their 6-byte FCS, to not exceed 2040 bits, as is required by the error-correcting code used with the RF physical and MAC protocol. This short packet size limits the RAM required for packet buffering, to facilitate implementation of gateway modules using small amounts of circuit board space and electrical power.

Parcel Format

All Oxinet2 parcels consist of not more than 236 bytes, including a 10-byte parcel (presentation layer) header, and a 0- to 226-byte information field.

Limiting the parcel size to 236 bytes permits the maximum-length parcel to fit into the information field of an Oxinet2 packet.

MAC Headers

Each Oxinet2 packet or parcel is encapsulated in the appropriate frame for the network type over which the packet or parcel is being transmitted. The MAC header occupies the first several bytes of the resulting MAC frame. The contents of these bytes are defined by the particular physical, MAC, (and LLC) layers in use. Oxinet2 session services do not directly utilize the physical and MAC header information in the MAC frames containing Oxinet2 packets.

The network driver in each station must supply the appropriate MAC and LLC header information for each outgoing frame. In some cases this information is supplied using values saved from previously received incoming frames.

Parcel Header

The parcel is the basic unit of information transfer between the application layer entities. Each parcel begins

with a 10-byte parcel header. The parcel header is used to: (1) direct the parcel to the appropriate application layer entity; (2) define the parcel's length; (3) designate any special properties of the parcel; and (4) identify the action of the parcel, thereby conveying a specific command or datum and implicitly defining the format of any data in the parcel's information field.

The parcel header is comprised of nine fields, the last two of which may have arbitrary contents in data parcels. The contents and usage of each of these fields are detailed in the following section.

OFFSET	NAME	USAGE
0	STAG	Source tag
1	DTAG	Destination tag
2	PILN	Parcel information Length
4	OID	Originator service identifier
5	RID	Recipient service identifier
6-7	ACT	Action code
8	ASEQ	Application sequence number
9	CCODE	Completion code

Source Tag (STAG)

The source tag field contains the tag value that identifies the instrument or workstation-based function selling this parcel. Specific tag values are assigned to each type of instrument and workstation-based function.

Destination Tag (DTAG)

The destination tag field contains the tag that identifies the intended recipient instrument or workstation-based function for this parcel. In the case of parcels that contain replies to previous requests, or parcels that are periodic data reports in response to previous reporting action, the DTAG of the reply parcel of data parcel contains the value obtained from the STAG of the corresponding request parcel.

Specific tag values are assigned to each type of instrument and workstation-based function.

Parcel Information Field Length (PILN)

The parcel information length field contains the total number of bytes in the information field of the parcel. The minimum information field length is 0 bytes, and the maximum length information field length is 226 bytes.

Parcel Flags (PCLF)

The parcel flags field contain several indicators relating to the handling of the parcel by the network control functionality. The usage of this field is independent of the ACT field, which relates to the handling of the parcel by the recipient at the application layer.

The usage of PCLF bits are detailed below:

BIT	NAME	USAGE
7-6	K (Kind)	These bits indicate the kind of action conveyed in this parcel. This indication is used to determine the length of the parcel header, and to implement reply timeouts; but it is not validated against the action code by parcel handling routines. These bits are encoded as follows: 00 = reserved 01 = request 10 = reply 11 = data
5	D (Destination type)	This bit indicates the type of addressing used with this parcel. If set = 0 the DTAG is matched against the instrument's ITP. If set = 1 the DTAG is matched against the station's

- (1) their own ITP (identifying a workstation-based function) in the STAG field;
- (2) the intended recipient's ITP in the DTAG field;
- (3) their own service ID in the OID field;
- (4) the intended recipient's service ID in the RID field; and
- (5) the D-bit in the PCLF field set to zero.

DID-Directed Parcels

Parcels sent between entities within a satellite cluster are generally addressed to services at specific stations. These parcels contain:

- (1) their own DID in the STAG field;
- (2) the intended recipient's DID in the DTAG field;
- (3) their own service ID in the OID field;
- (4) the intended recipient's service ID in the RID field; and
- (5) the D-bit in the PCLF field set to one.

Network Protocol

The network protocol is used to transfer Oxinet2 packets between workstations and satellite clusters over ARCNET and the Oxinet2 RF network. An ARCNET segment is used for all communication into and out of workstations. An ARCNET segment is used for all communication into and out of workstations. Satellite clusters may be (1) directly connected to the ARCNET segment, using an ARCNET gateway module, or (2) indirectly connected to the ARCNET segment using an RF gateway module communicating through an RF bridge on the ARCNET segment.

Connections to a network segment take place through gateway modules. While multiple gateways may provide simultaneous physical attachment to one or more network segments, no more than one gateway may be engaged in active session communication with a workstation at any time.

The parcel is the basic unit of information transferred between the application layer entities. One or more parcels may be carried in each Oxinet 2 packet; however, parcels are never split across packet boundaries. Parcels communicated through any particular session are delivered according to the parcel precedence rules defined below.

Network MAC and LLC Usage

Oxinet2 packets are encapsulated into ARCNET packets or RF packets as appropriate for the link layer facilities being used on the particular network segment.

Usage with ARCNET

When Oxinet2 packets are sent via ARCNET, the necessary MAC information is the (1) frame type (FT), which is a code of 01h, indicating data packet, provided by the station's network controller chip; (2) source node ID (SID), provided by the station's network controller chip; (3) destination node ID (DID), provided by the network driver based on information associated with the DHD of the session, and transmitted twice by the network controller chip as a means of ensuring correct MAC layer address validation at the recipient station; and (4) continuation pointer (CP), that encodes the length of the encapsulated Oxinet2 packet plus the LLC header byte.

In addition to the MAC header, ARCNET uses a MAC trailer containing a 2-byte frame check sequence (FCS), calculated by the station's network controller chip. This frame check sequence, calculated using the CRC-16 polynomial, is present to permit rejection of frames received with data errors. The ARCNET FCS does not provide error correction, which is deemed unnecessary due to a bit error rate on ARCNET which is less than $1 \text{ in } 10^{12}$.

Usage with the Oxinet2 RF Network

When Oxinet2 packets are sent via the RF network, the spread-spectrum coding used by the RF transceivers pro-

vides substantial rejection of non-Oxinet2 signals. The encapsulation includes a 4-byte header containing a starting delimiter, frame type (repeated), and frame length; and a 7-byte trailer containing a Reed-Solomon error correcting code and an ending delimiter. The error correcting code permits correction of burst errors up to two byte sin length, and is expected to permit over 98% of RF packets to be used without re-transmission despite a bit error rate which could be as large as $1 \text{ in } 10^6$.

MAC Address Usage Rules

MAC frames on ARCNET and the RF network may be sent with destination addresses that contain either (1) zero, indicating a broadcast to all stations; or (2) a non-zero value, indicating that the frame is intended to be received by a single, designated station.

On the RF network, the broadcast destination is only allowed on packets sent from satellite stations to RF bridges. In all cases, the source addresses of MAC frames contain the assigned MAC address of the sending station. The destination address of a MAC frame being sent as a reply should contain the value obtained from the source address of the frame to which the outgoing frame is a reply.

LLC Header

The sole purpose of the LLC layer in Oxinet2 packets is to permit multiple protocol types to coexist, concurrently on the same MAC layer. This is only significant with ARCNET, since the Oxinet2 RF network is proprietary, and the spread-spectrum coding used as the physical layer of the RF network provides substantial rejection of non-Oxinet2 signals within the same radio frequency band.

The first byte after the MAC header on ARCNET is a protocol type identifier byte, which constitutes a one-byte LLC header. No LLC header is used on the Oxinet2 RF network.

Special MAC and LLC Usage Rules

Upon receipt of a MAC frame from ARCNET the protocol type identifier is checked, and only frames with the Oxinet2 system code are processed. Frames with the protocol types are passed to the protocol handler for that protocol type (if present) or are discarded if there is no protocol handler available for that protocol type.

The MAC layer length information is used to determine the total number of bytes in the frame, which may exceed the number of bytes in the Oxinet2 packet. This situation is detected when the MAC information length exceeds the Oxinet2 packet information length. In such cases the extra bytes at the end of the frame (beyond the Oxinet2 information length) are not used.

This length difference may occur because the ARCNET network controller chip requires the MAC information of a data packet to be justified to the end of its data buffer. In certain implementations, especially those based on limited-performance microcontrollers in gateway modules, implementers may choose to pre-allocate a MAC information field of a particular length, even if the Oxinet2 packet is shorter, to avoid the overhead of repeated, memory-to-memory movement of overlapping areas in the packet buffer RAM.

Upon receipt of an Oxinet2 packet from either ARCNET or the Oxinet2 RF network while no session is in progress an incoming packet is only accepted if directed to session zero; and

the only packets accepted are those containing parcels with the following network control actions:

- CONN_REQ, to establish a session,
- RDEE_REQ, to read EEPROM parameters,
- STAT_REQ, to read out status and statistics, and
- STST_REQ, to read out statistics.

such detection and elimination at the presentation layer or application layer.

The possibility of redundant reception of packets exists because the MAC layer acknowledgement, which serves as positive acknowledgement of packet reception, is transmitted in a separate MAC frame from the data packet being acknowledged. On both ARCNET and the Oxinet2 RF network, it is possible for a packet to be successfully received by the target, then for the MAC layer positive acknowledgement to be incorrectly received (or never to reach) the initiator. In such a case, the initiator will automatically retransmit the packet, causing a redundant reception of the same packet by the target.

Each transmitter state machine maintains a value for use in generating the TXSEQ field in the Oxinet2 headers of outgoing packets. This value is incremented by one upon completion of the packing of each packet for transmission. Once set, the TXSEQ byte of any packet remains unchanged, even if the packet needs to be retransmitted.

Each receiver state machine retains the value of the TXSEQ field from the last validly received packet. Upon successful receipt of any packet during a session, the receiver state machine discards the packet if the low-order bit of the TXSEQ field is equal to the low-order bit of the retained value from the TXSEQ of the previous valid packet of the same session. When no session is in progress the TXSEQ field is ignored by the receiver state machine.

The sole purpose of transmit sequencing in the Oxinet2 packet header is to permit redundantly received packets (within a session) to be discarded at the transport layer. The transmitter sequence number is not used to detect missing packets.

Packet Acknowledgement

The transmitter state machine uses the MAC layer acknowledgement facilities of the underlying network segment to determine whether packets have reached their destinations successfully. The acknowledgement rules are as follows: the receipt of a MAC layer positive acknowledgement by the initiator of any packet indicates successful delivery of that packet to the target on this network segment. Upon receipt of the positive acknowledgement, the transmitter state machine at the initiator may reclaim the transmit buffer used for the delivered packet and prepare to send a new packet during the next transmission opportunity.

On ARCNET receipt of an ACK frame constitutes positive acknowledgement to packet delivery. This is detected as both TA=1 and TMA=1 in the status register of the ARCNET controller chip. On the Oxinet2 RF network, receipt of a positive acknowledgement frame or receipt of a data packet frame with a frame type of 3 Ch constitutes positive acknowledgement of packet delivery.

The receipt of a MAC layer negative acknowledgement by the initiator of any packet, or receipt of no acknowledgement during a defined acknowledgement interval, indicates that the packet has not successfully reached the target on this network segment. The transmitter state machine must retransmit such non-delivered packets until either a positive acknowledgement is received or the session timeout occurs.

On ARCNET the lack of an ACK frame within 78 us of the end of transmitting a data packet frame constitutes negative acknowledgement to packet delivery. This is detected as TA=1 and TMA=0 in the status register of the ARCNET controller chip. Also, the inability to transmit a packet due to continuous failures of free buffer inquiries is considered to constitute negative acknowledgement. This condition, known as a "blocked receiver", is detected when the Excessive NAK bit in the ARCNET controller chip is set,

indicating at least 128 consecutive free buffer inquiry failures. An EEPROM parameter, with a default value of 1, is used to define the number of successive Excessive NAK indications constitute a blocked receiver condition.

Due to the typical reasons for a TA=1, TMA=0 condition on ARCNET, the recommended practice for ARCNET packet driver implementation is to retry an unsuccessfully delivered outgoing transmission one time before indicating a negative acknowledgement condition. On the Oxinet2 RF network, receipt of a negative acknowledgement frame or receipt of a data packet frame with a frame type of C3h constitutes negative acknowledgement of packet delivery. Also, no response of any type for a full synchronization interval following transmission of a data packet frame from an RF bridge, or no response of any type during the outbound interval following transmission of a data packet frame to an RF bridge, is treated as being equivalent to negative acknowledgement.

In cases where the target of a directed packet on the RF network is an RF bridge, that bridge must only acknowledge receipt of the packet if there is an active route for that packet in the bridge's routing table, and there are no more packets awaiting transfer via the same route as there are slots assigned for that route in the bridge's synchronization period. Once an RF bridge has positively acknowledged a received RF packet, the bridge must deliver that packet to the designated destination (on the ARCNET) unless prevented from doing so by failure of the network or failure of the recipient station.

Parcel Handling

Application layer entities communicate using parcels. Individual parcels are provided to the session layer for transfer to the session partner. One or more of these parcels are packed into Oxinet2 packets for transfer over the network segment(s) connecting the clusters on which the session partners reside.

The MAC layers used on the ARCNET and Oxinet2 RF network provide reliable confirmation of packet delivery. However, there may be conditions, caused either by physical interference on the network medium or by exhaustion of buffer space at the sending or receiving station, under which outgoing packets are undeliverable, or are deliverable at a lower rate than needed to keep up with the delivery of outgoing parcels from the presentation layer interface.

Parcel Precedence

In order to ensure that the most important parcels are delivered in a timely manner whenever any communication is possible; and to ensure that, if undelivered parcels need to be discarded the least important are discarded first, a system of parcel precedence is used. There are four levels of precedence, listed below in order of decreasing urgency and increasing discardability:

Precedence 0 is used for reporting alarms and status for which delivery is mandatory if any communication is possible. These parcels require timely delivery, so if a newer precedence 0 parcel is presented for delivery, any older parcel from the same source (identified by OID) awaiting delivery is discarded. The newer parcel replaces the older parcels to ensure reporting of the most recent information.

Precedence 1 is used for request and reply parcels. These parcels are unlikely to cause buffer exhaustion because there are generally far more buffers available than there can be meaningful outstanding requests. Precedence 1 parcels are only discarded when insufficient buffer space is available to hold all precedence 0 and 1 parcels.

Precedence 2 is used for transfer of "non-snapshot" waveform data, which is discardable only if absolutely

received connect request with a CONN_REQ, indicating either success or failure of the session establishment. The CONN_REQ parcel is always returned in a directed packet to the station designated by the SHD of the packet that contained the CONN_REQ.

Connection Requests

CONN_REQ parcels are sent in broadcast packets in order to reach all active session administrators, even if the satellite cluster and the workstation are attached to different network segments. This permits session establishment to occur without requiring the satellite clusters to have knowledge of the network topology.

The CONN-REQ MAC, packet, and parcel headers are used as follows:

The MAC header (ARCNET only) contains

- (1) the SID of the gateway requesting the connection and
- (2) a broadcast DID.

The Oxinet2 packet header contains

- (1) an SHD with the STID of the gateway and SSN=1,
- (2) a DHD of zero, meaning broadcast (including SSN=0 so the request goes to session zero), and
- (3) a PCT of one (for the CONN_REQ parcel).

The CONN_REQ parcel header contains

- (1) an STAG with the ITP of the instrument in which the active gateway is installed,
- (2) a DTAG of EOh, indicating network control functionality of the workstation,
- (3) an OID of 02h indicating session manager service,
- (4) an RID of FOH, indicating session administrator service, and
- (5) an ACT of CONN_REQ.

In cases where there are multiple gateways in a satellite cluster, inactive gateways that become able to communicate on a network segment may attempt to establish a session with a workstation in preparation for an attempt to force a session manager handoff. This procedure uses a make-before-break discipline, wherein

- (1) the alternate gateway sends a CONN_REQ, with OID of 40h or 44h and a non-zero value in the cur_ssn field of the parcel;
- (2) if a successful CONN_REQ is received, the gateway attempts to become session manager by sending a HDOF_REQ to the session manager;
- (3) if the handoff is successful, the alternate gateway becomes session manager and sends a DISC_REQ to the previous session manager to be forwarded to the workstation via the old session; and
- (4) if the handoff is unsuccessful, the alternate gateway goes inactive, sending a DISC_REQ to the workstation to terminate the new session.

Connection Replies

CONN_REQ parcels are sent in directed addressed to the sender of the CONN_REQ parcel. In the case of replies to requests forwarded by RF bridges, the bridge updates its routing table using address information from the reply, as further defined under "Routing" below.

The CONN_REQ MAC, packet, and parcel headers are used as follows:

- (1) the SID of the workstation generating the CONN_REQ, and
 - (2) the DID copied from the SID of the CONN_REQ.
- The Oxinet2 packet header contains
- (1) an SHD with the STID of the workstation and an SSN value assigned by the session administrator,

- (2) a DHD copied from the SHD of the CONN_REQ, and
 - (3) a PCT of one (for the CONN_REQ parcel).
- The CONN_REQ parcel header contains

- (1) an STAG of EOh, indicating network control functionality of a workstation,
- (2) a DTAG copied from the STAG of the CONN_REQ,
- (3) An OID of FOH, indicating session administrator service,
- (4) an RID copied from the OID of the CONN_REQ,
- (5) an ACT of CONN_REQ,
- (6) an ASEQ copied from the ASEQ of the CONN_REQ, and
- (7) a CCODE indicating the success or failure of the session establishment attempt.

Reply Processing

Because more than one workstation may be active on the network, it is possible for a single CONN-REQ to result in more than one CONN_REQ with successful completion status. The priorities for acceptance of connect requests by session administrators, and for acceptance of connect replies by satellite session managers are

- (1) the workstation responsible for this administrative domain,
- (2) the workstation that was the partner in the last session established from this satellite, or
- (3) any other workstation that will adopt this satellite.

The STID of the last session partner is retained by each satellite as an EEPROM parameter, and the value of this parameter is included in subsequent CONN_REQ parcels, to facilitate implementation of these priorities. To minimize the number of EEPROM write cycles used to perform this function, a satellite waits for at least 150% of the session timeout EEPROM parameter, after a successful CONN_REQ from a workstation other than the last custodian, to determine whether a successful CONN_REQ from the last custodian is also pending. This is in addition to the (standard) read-before-write used on all EEPROM updates.

A satellite attempting to establish a session typically accepts the first successful CONN_REQ that is received. Then, if a subsequent CONN_REQ from a higher priority session partner (as listed above) is received, disconnects from the first partner to accept the second partner. The DISC_REQ indicates the reason for this disconnection as "better partner".

If the subsequent CONN_REQ is from a lower priority session partner, the subsequent reply is discarded. By discarding redundant or unsuitable replies, the unnecessary session will terminate due to inactivity at the end of a single session timeout interval.

Because the CONN_REQ is sent as a broadcast, there is no direct acknowledgement of receipt. A station sending a CONN_REQ should wait up to the session timeout interval for a reply. Short-duration retries may be made for up to 3 times the session timeout interval, after which the periods between retries must be increased substantially.

In the case of satellites that are unable to establish a session via a wired network segment, periodic retries do not constitute a threat of generating excessive network traffic nor will they deplete the batteries of the satellite station. However, in the case of satellites attempting to establish a session via an RF network segment, both the transmission bandwidth and the battery power consumed for these periodic retries can be a problem. Accordingly, RF bridges indicate, once per second, using facilities of the RF MAC protocol, the current count of workstations for which that

(2) Waiting for Transmission Opportunity (WT_TXOP), which is used after a set of outgoing parcels has been packed into a packet and is ready for transmission or retransmission; and

(3) Waiting for Acknowledgement (WT_ACK), which is used after each packet transmission attempt while waiting for a positive or negative acknowledgement.

One integer value is maintained by the transmitter state machine TXSEQ, which represents the current sequence number used for transmission of packets, and is used to allow the receiver to detect and discard redundantly received packets.

One EEPROM parameter is used in the operation of the transmitter state machine. The Session Timeout (Ts) defines the interval following packet transmission during which positive acknowledgement, or an incoming packet, must be received from the session partner. The default value for the session timeout is 7 seconds.

Events

The events that can cause state transitions in the transmitter state machine include:

"Parcels Available", which indicates that one or more parcels are enqueued,

"near TXOP", which indicates that a transmission opportunity will occur by the time the available parcels are packed for transmission.

For communication over ARCNET, the near TXOP event is always asserted, since transmission opportunities on ARCNET occur no less frequently than once every 50 ms. For communication over the Oxinet2 RF network, the near TXOP event is asserted a sufficient interval prior to this station's assigned slot for the station's control processor to pack and otherwise prepare a full-length outgoing packet for transmission. This "early warning" of an impending transmission opportunity permits efficient usage of the parcel precedence mechanism on a network where transmission opportunities typically occur only once every second.

"TXOP", which indicates the occurrence of a transmission opportunity on the underlying network.

For communication over ARCNET, the TXOP event is always asserted, since arrival of the ARCNET token at a given station cannot be accurately predicted nor detected, and because the act of "transmitting" a packet on ARCNET, as seen by the low-level packet driver, involves enabling the transmitter on the ARCNET controller chip.

For communication over the Oxinet2 RF network, the TXOP event is asserted at the start of the station's assigned slot.

"Positive ACK", which indicates the receipt of a positive acknowledgement to the most recent packet transmission.

"Negative ACK", which indicates the receipt of a negative acknowledgement to the most recent packet transmission.

"Ts Timeout", which indicates the expiration of the session timeout interval.

Transmitter State Transition Diagram

Operation of the transmitter state machine is characterized by the diagram of FIG. 8. A narrative description of operation of this state machine follows:

Upon initialization of a transmitter state machine, or upon completion of processing for any packet, state WT_PAR is entered. In state WT_PAR the transmitter

state machine is idle, awaiting the coincident availability of an impending transmission opportunity and one or more parcels to transmit. If outgoing parcels have been enqueued during the processing of the previous packet, only the near TXOP event (which is always asserted when using ARCNET) is needed to exit WT_PAR state.

Upon exit from WT_PAR state a packet is packed using as many of the available parcels will fit into the packet (or all that are on the queue if the packet is not full). Parcels are taken from the outgoing parcel queue in order of their precedence, and within each precedence level in order from oldest to newest. After the packet is prepared for transmission, the session time (Ts) is reset (and restarted), and state WT_TXOP is entered.

In state WT_TXOP the transmission state machine is awaiting the arrival of a transmission opportunity. On ARCNET this state is effectively null, since the continuous assertion of the TXOP event causes immediate exit from this state.

Upon exit from WT_TXOP state the packet is transmitted, and state WT_ACK is entered.

On ARCNET, "transmitting" the packet refers to placing the packet into the packet buffer of the ARCNET controller chip and issuing an enable transmitter command to that chip.

In state WT_ACK the transmitter state machine is awaiting an acknowledgement to the previously transmitted packet or a Ts timeout. Receipt of a positive acknowledgement indicates successful delivery of the packet, in which case (1) the Ts timer is reset, (2) the outgoing packet's buffer is freed, (3) the value of TXSEQ is incremented by 1, and (4) state WT_PAR is entered.

Receipt of a negative acknowledgement indicates unsuccessful delivery of the packet, in which case state WT_TXOP is entered to cause retransmission of the packet. On stations that support multiple sessions (workstations and RF Bridges), no more than two successive transmission opportunities are to be used to attempt delivery of any given packet before scanning the queue of outgoing packets and attempting delivery of any packets pending from other sessions.

Occurrence of a Ts timeout (while in WT_PAR or WT_ACK states) causes (1) the session in progress (if any) to be terminated, (2) the presentation layer to be informed of the timeout condition, (3) the outgoing packet's buffer to be freed, (4) the value of TXSEQ to be incremented by 1, and (5) state WT_PAR to be entered.

Receiver Control State Machine

The receiver state machine shown in FIG. 9 handles all incoming packets for a given partner; validates the Oxinet2 header on received packets; and inhibits the receive function when no free receive buffers are available for additional incoming packets.

Receiver states are

Waiting for Packet (WT_PKT), which is used while waiting for an incoming packet to be received; and

Waiting for Buffer (WT_BUF), which is used when all receive buffers are occupied and no further reception is possible until at least one buffer is freed by higher-layer functionality.

One integer value is maintained by the receiver state machine. RXSEQ, which holds the value from the TXSEQ field of the last validly received packet. The low-order of this variable must be the complement of the low-order bit of the TXSEQ field in the next incoming packet for the new packet to be accepted as valid.

ECHO_REQ Parcel Info Field:		
OFFSET	FIELD	USAGE
1-225		Used to hold timestamp reports

Connect

The connect action establishes a session if the requester is an appropriate partner for the recipient and the recipient has the capacity to create an additional session at the time the request is received. CONN_REQ parcels are sent by session manager services at satellite clusters to session administrator services at workstations. The SSN field in the SHD of the packet that carries the CONN_REQ parcel should be=1 and the SSN field in the DHD should be=0. The session administrator at the workstation assigns the SSN for the satellite to use in the DHD of subsequent packets directed to the workstation during this session, and returns the value to the requester in the SHD of the packet used to carry the CONN_REQ parcel.

If a satellite cluster does not receive a CONN_REQ within the session timeout interval, or the CONN_REQ that is received contains a completion code other than zero (successful) or a rejection code indicating not to retry, the CONN_REQ is repeated up to a total of three times in rapid succession. If a connection has still not been successfully established, the station repeats the CONN_REQ actions once per minute until a session is successfully established. This one minute retry cycle can be overridden by various local actions, such as the activation of an instrument-specific control sequence to cause immediate retry of the connection attempt.

Several parameters about the device sending the CONN_REQ are included in the request parcel for use at the destination in determining whether the request should be granted. If these parameters indicate that the session should not be established (due to incompatible protocol versions, etc.) the request is rejected with the appropriate completion code.

CONN_REQ parcels are never ignored. If the CONN_REQ is received, a CONN_REQ is generated whether or not the request is successful. If the completion code field contains any value other than 0, the session has not been successfully established, and the code specifies the reason for failure.

CONN_REQ Parcel Info Field:		
OFFSET	FIELD	USAGE
0-3	iid	Instrument ID
4	hw_rev	Hardware revision level
5	sw_rev	Software revision level
6	min_prot	Minimum supportable protocol version
7	max_prot	Maximum supportable protocol version
8	cur_ssn	Current SSM (= 0 if unconnected)
9-11	last_stid	STID of last (or current) session partner
12-15	last_conn	Time last connected
16-19	time	Current time
20-22	location	Last known location
23-26	last_loc	Time location last received
27	ee_slrc	SLRC of EEPROM

CONN_REQ Parcel Info Field:		
OFFSET	FIELD	USAGE
0	reason	Reason for successful connection
1-3	stid	Replying station's STID
4	hw_rev	Hardware revision level
5	sw_rev	Software revision level
6	prot_ver	Protocol version to use

When the CCODE in the CONN_REQ parcel header indicates successful completion (zero), byte 0 of the INFO field contains bit-encoded information on the reason for granting this request:

- bit 0 is=1 if the request has been granted due to an open adoption (not currently used);
- bit 1 is=1 if the replying workstation is this requester's last custodian;
- bit 2 is=1 if the requester is in the replying workstation's administrative domain; and
- bits 3-7 are reserved.

This bit encoding permits CONN_REQ parcels received from partners with higher connection priority to be determined by a numerically greater reason value.

Sessions may only be established between partners that can both support a common protocol version. The specific protocol version to use is defined in the prot_ver byte of the CONN_REQ parcel.

CONNECT Completion Codes:	
VALUE	USAGE
00h	Successful completion of CONNECT action
04h	Action rejected, no more sessions available at workstation
11h	Unsupported protocol version(s)
12h	Invalid adoption
13h	Incorrect administrative domain
84h	CONNECT action rejected, do not call back

Disconnect

The disconnect action terminates a session. This action is rejected if no session is in progress. Either session partner can initiate a DISC_REQ. The DISC_REQ is the last parcel sent over the session being terminated. If a workstation receives a DISC_REQ with the B-bit in the PKTF byte set=1, the workstation must expunge the route used for the session being disconnected by sending a UNRT_REQ to the appropriate RF bridge.

DISC_REQ Parcel Info Field:		
OFFSET	FIELD	USAGE
0	reason	Reason for disconnection
1-3	stid	Sending station's STID
4-6	new_stid	STID of new session manager (only valid if handoff)

Byte 0 of the INFO field contains bit-encoded information on the reason(s) the session is being disconnected:

- bit 0 is=1 if the requester has received a CONN_REQ from a more appropriate session partner;
- bit 1 is=1 if the requesting cluster has obtained a connection that uses a more appropriate source of power;
- bit 2 is=1 if the requesting cluster has obtained a connection via a more appropriate network medium;
- bit 3 is=1 if the requester has suffered a loss of power or low-battery condition;

-continued

STAT_REP Parcel Info Field:		
OFFSET	FIELD	USAGE
3-6	rsiid	Replying station's IID
7-10	timest	Current timestamp
11-14	sv_timest	Saved timestamp
15	post	Power-on self-test status (0 = no errors)
16	cum_post	Logical-OR of POST results from Gw or SM
17	ec_slrc	SLRC of EEPROM data
18	connst	Connection status (0 = not in session)
19	gatest	Gateway status (0 = network not available)
20	locst	Location status (0 = not located)
21-23	irloc	Last known IR location
24-27	loc_time	Timestamp when location last received
28	inst_ct	Number of instruments in table (only relevant from session manager)

NOTE:

All counters and saved timestamps are reset to zero at power-on reset. The current timestamp is set at power-on reset from the persistent time sense facility, if present, or is set to zero if no time sense is available.

STATUS Completion Codes:	
VALUE	USAGE
00h	Successful completion of STATUS action

Module Configuration

The module configuration action is used by the instrument gateway during instrument initialization to determine the necessary information to build the local service table and initialize the trend buffers.

MCON_REQ Parcel Info Field:		
OFFSET	FIELD	USAGE
{INFO field is not used}		

MCON_REP Parcel Info Field:		
OFFSET	FIELD	USAGE
0-2	rsiid	Replying station's STID
3-6	rsiid	Replying station's IID
7	post	Power-on self-test status (0 = no errors)
8	gwstat	Gateway status (see HDOF_REQ)
9	trends	Number of trend bytes reported each 10 sec.
10	tr_src	Type of time reference source
11	tr_dst	Type of time reference destination
12	services	Number of local services (1-15)
13-27		List of local service names

MODULE CONFIGURATION Completion Codes:	
VALUE	USAGE
00h	Successful completion of MODULE CONFIGURATION action

Statistics

The statistics action is used to determine various operation statistics pertaining to the repier.

STST_REQ Parcel Info Field:		
OFFSET	FIELD	USAGE
{INFO field is not used}		

STST-REP Parcel Info Field:		
OFFSET	FIELD	USAGE
0-2	rsiid	Replying station's STID
3-6	rsiid	Replying station's IID
7-10	timest	Current timestamp
11-14	sv_timest	Saved timestamp
15-18	tx_pkt	Number of parcels transmitted
19-22	tx_pkt	Number of packets transmitted
23-26	tx_byte	Number of bytes transmitted
27-30	null_dst	Number of nonexistent destinations
31-34	blk_rcv	Number of blocked receivers
35-38	rx_pkt	Number of parcels received
39-42	rx_pkt	Number of packets received
43-46	rx_byte	Number of bytes received
47-50	disc_pkt	Number of parcels discarded
51-54	disc_pkt	Number of packets discarded
55-58	disc_byte	Number of bytes discarded
59-62	bad_fr	Number of invalid frames
63-66	recons	Total number of reconfigurations
67-70	my_recons	Number of reconfigurations from the station
71-74	corr	Number of packets with ECC errors corrected
75-78	uncorr	Number of packets with uncorrectable ECC errors

STATISTICS Completion Codes:	
VALUE	USAGE
00h	Successful completion of STATISTICS action

As will be understood by those familiar with the art, the present invention may be embodied in other specific forms without departing from the spirit or essential characteristics thereof. For example, the number and location of the bytes in the headers could be varied. Accordingly, the disclosure of the preferred embodiment of the invention is intended to be illustrative, but not limiting, of the scope of the invention, which will be set forth in the following claims.

What is claimed is:

1. A network having a plurality of nodes coupled together over a communication medium, comprising:
 - at least first and second nodes coupled to said communication medium;
 - a plurality of virtual services associated with at least one device coupled to said first node;
 - means, associated with said virtual services, for creating parcels containing an information field and a virtual destination service header;
 - service table means, at said first node, for storing a physical address corresponding to a virtual destination service address;
 - means, at said first node, for transmitting said parcels over said communication medium to said physical address;
 - means, at said second node, for receiving said parcels from said communications medium; and
 - means, at said second node, for directing each of said parcels to the virtual destination service specified in said virtual destination service header;

45

a source tag byte with a first range of values identifying instruments and a second range of values identifying workstation-based functions;

a destination tag byte with said first range of values identifying instruments and said second range of values identifying said workstation-based functions;

a parcel information length byte indicating the length of said parcel information field;

a parcel flag byte for indicating a kind of action conveyed in said parcel information field, a type of addressing used and a precedence level;

an originator service identifier byte indicating a source virtual service;

a recipient service identifier byte indicating a destination virtual service; and

a pair of action code bytes indicating an action to be taken in processing said parcel information field at an application layer in said destination virtual service.

5. A gateway for coupling an instrument to a network comprising:

means for receiving and temporarily storing a plurality of parcels from at least one virtual service in said instrument, each parcel containing a parcel information field and a virtual address header;

means for multiplexing said plurality of parcels into a single packet, said single packet having a packet header and a packet information field, said packet information field containing said plurality of parcels;

means for transmitting said single packet over said network during a single session; and

means for constructing said packet header to comprise:

four source handle bytes including three station ID bytes identifying said gateway and a first session number byte;

four destination handle bytes including three station ID bytes identifying a destination gateway and a second session number byte;

a transmitter sequence number byte identifying said packet;

a packet flag byte indicating whether there have been any errors or any reconfiguration;

an information length byte indicating the length of said packet information field; and

a parcel count byte indicating the number of parcels in said packet information field.

6. A gateway for coupling an instrument to a network comprising:

means for receiving and temporarily storing a plurality of parcels from at least one virtual service in said instrument, each parcel containing a parcel information field and a virtual address header;

means for multiplexing said plurality of parcels into a single packet, said single packet having a packet header and a packet information field, said packet information field containing said plurality of parcels;

means for transmitting said single packet over said network during a single session; and

means for constructing said virtual address header to comprise:

a source tag byte with a first range of values identifying instruments and a second range of values identifying workstation-based functions;

a destination tag byte with said first range of values identifying instruments and said second range of values identifying said portion of a computer connected to said first node;

46

a parcel information length byte indicating the length of said parcel information field;

a parcel flag byte for indicating a kind of action conveyed in said parcel information field, a type of addressing used and a precedence level;

an originator service identifier byte indicating a source virtual service;

a recipient service identifier byte indicating a destination virtual service; and

a pair of action code bytes indicating an action to be taken in processing said parcel information field at an application layer in said destination virtual service;

wherein said virtual address header contains presentation layer header information and said packet header contains network, transport and session address headers.

7. A network having a plurality of nodes coupled together over a communication medium, comprising:

at least first and second nodes coupled to said communication medium;

a plurality of virtual services associated with at least one device coupled to said first node;

means, at said first node, for receiving and temporarily storing a plurality of parcels from said virtual services, each parcel containing a parcel information field and a virtual address header;

means, at said first node, for multiplexing said plurality of parcels into a single packet, said single packet comprising a packet header and a packet information field, with said plurality of parcels being in said packet information field of said single packet;

means, at said first node, for transmitting said single packet over said communication medium during a single session;

means, at said second node, for receiving said single packet from said communications medium;

means, at said second node, for demultiplexing said received packet into a plurality of parcels; and

means, at said second node, for directing each of said parcels to the virtual address specified in said virtual address header information;

wherein said virtual address header comprises:

a source tag byte with a first range of values identifying instruments and a second range of values identifying workstation-based functions;

a destination tag byte with said first range of values identifying instruments and said second range of values identifying said workstation-based functions;

a parcel information length byte indicating the length of said parcel information field;

a parcel flag byte for indicating a kind of action conveyed in said parcel information field, a type of addressing used and a precedence level;

an originator service identifier byte indicating a source virtual service;

a recipient service identifier byte indicating a destination virtual service; and

a pair of action code bytes indicating an action to be taken in processing said parcel information field at an application layer in said destination virtual service.

8. A network having a plurality of nodes coupled together over a communication medium, comprising:

at least first and second nodes coupled to said communication medium;



US006213942B1

(12) **United States Patent**
Flach et al.

(10) Patent No.: **US 6,213,942 B1**
 (45) Date of Patent: **Apr. 10, 2001**

(54) **TELEMETER DESIGN AND DATA
 TRANSFER METHODS FOR MEDICAL
 TELEMETRY SYSTEM**

(75) Inventors: **Terry E. Flach, Altadena; Michael D.
 Stoop, Aliso Viejo, both of CA (US)**

(73) Assignee: **Vitalcom, Inc., Tustin, CA (US)**

(*) Notice: Subject to any disclaimer, the term of this
 patent is extended or adjusted under 35
 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/315,254**

(22) Filed: **May 20, 1999**

Related U.S. Application Data

(62) Division of application No. 08/675,594, filed on Jul. 3, 1996,
 now Pat. No. 5,944,659.

(60) Provisional application No. 60/006,600, filed on Nov. 13,
 1995.

(51) Int. Cl.⁷ **A61B 5/00; A61F 2/02**

(52) U.S. Cl. **600/300; 600/301; 128/903;
 128/904**

(58) Field of Search **600/300, 301,
 600/481, 345, 347, 500, 529, 544-545,
 595; 128/900, 903, 904, 905; 705/2, 3**

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,603,881 9/1971 Thornton .
 3,826,868 7/1974 Nugent .
 3,925,762 12/1975 Heitlinger et al. .
 4,051,522 9/1977 Healy et al. .

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

06/02459 6/1994 (EP) .
 07/10465 5/1996 (EP) .
 2258960 2/1993 (GB) .
 2271691 4/1994 (GB) .

OTHER PUBLICATIONS

Product brochure titled "Wireless Connectivity by Pacific
 Communications, Inc.", 1993.

International Search Report, dated Oct.

Primary Examiner—Cary O'Connor

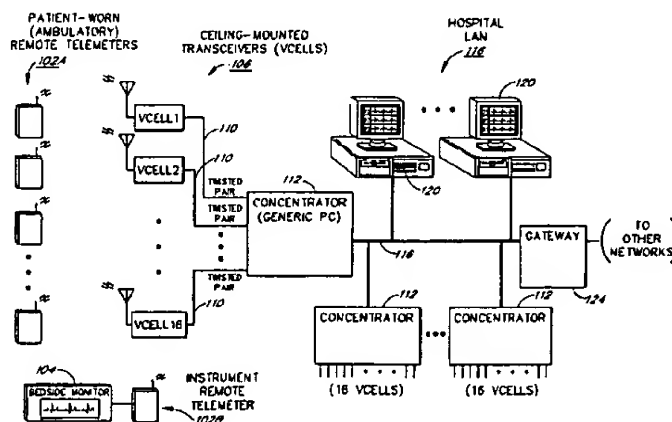
Assistant Examiner—Michael Astorino

(74) *Attorney, Agent, or Firm*—Knobbe, Martens, Olson &
 Bear, LLP.

(57) **ABSTRACT**

A medical telemetry system is provided for collecting the real-time physiologic data of patients (including ambulatory patients) of a medical facility, and for transferring the data via RF to a real-time data distribution network for monitoring and display. The system includes battery-powered remote telemeters which attach to respective patients, and which collect and transmit (in data packets) the physiologic data of the patients. The remote telemeters communicate bi-directionally with a number of ceiling-mounted RF transceivers, referred to as "VCELLs," using a wireless TDMA protocol. The VCELLs, which are hardwire-connected to a LAN, forward the data packets received from the telemeters to patient monitoring stations on the LAN. The VCELLs are distributed throughout the medical facility such that different VCELLs provide coverage for different patient areas. As part of the wireless TDMA protocol, the remote telemeters continuously assess the quality of the RF links offered by different nearby VCELLs (by scanning the frequencies on which different VCELLs operate), and connect to those VCELLs which offer the best link conditions. To provide a high degree of protection against multi-path interference, each remote telemetry maintains connections with two different VCELLs at a time, and transmits all data packets (on different frequencies and during different timeslots) to both VCELLs; the system thereby provides space, time and frequency diversity on wireless data packet transfers from the telemeters. The telemeters and VCELLs also implement a patient location protocol for enabling the monitoring of the locations of individual patients. The architecture can accommodate a large number of patients (e.g., 500 or more) while operating within the transmission power limits of the VHF medical telemetry band.

22 Claims, 11 Drawing Sheets





US 6,213,942 B1

Page 2

U.S. PATENT DOCUMENTS

4,539,710	9/1985	Dinsmore .	5,381,798	1/1995	Burrows .
4,958,645	9/1990	Cadell et al. .	5,416,695	5/1995	Stutman et al. .
5,153,584	10/1992	Engira .	5,441,047	8/1995	David et al. .
5,205,294	4/1993	Flach et al. .	5,458,122	10/1995	Hethuin .
5,238,001	8/1993	Gallant et al. .	5,458,123	10/1995	Unger .
5,270,811	12/1993	Ishibashi et al. .	5,502,726	3/1996	Fischer .
5,305,202	4/1994	Gallant et al. .	5,507,035	4/1996	Bantz et al. .
5,305,353	4/1994	Weerackody .	5,572,517	11/1996	Safadi .
5,309,920	5/1994	Gallant et al. .	5,579,001	11/1996	Dempsey et al. .
5,319,363	6/1994	Welch et al. .	5,579,378	11/1996	Arlinghaus, Jr. .
5,359,641	10/1994	Schull et al. .	5,579,775	12/1996	Dempsey et al. .
5,375,604	12/1994	Kelly et al. .	5,614,914	3/1997	Bolgiano et al. .



*FIG. 2*





FIG. 4




FIG. 5A




FIG. 5B




FIG. 6




FIG. 7




FIG. 8

TIMESLOT STATUS TABLE

900

ASSIGNED
FRAMES SINCE
LAST PACKET

SLOT 1	SLOT 2	SLOT 3	SLOT 4	SLOT 5	SLOT 6
✓	✓		✓	✓	
0	0	(NA)	3	0	(NA)

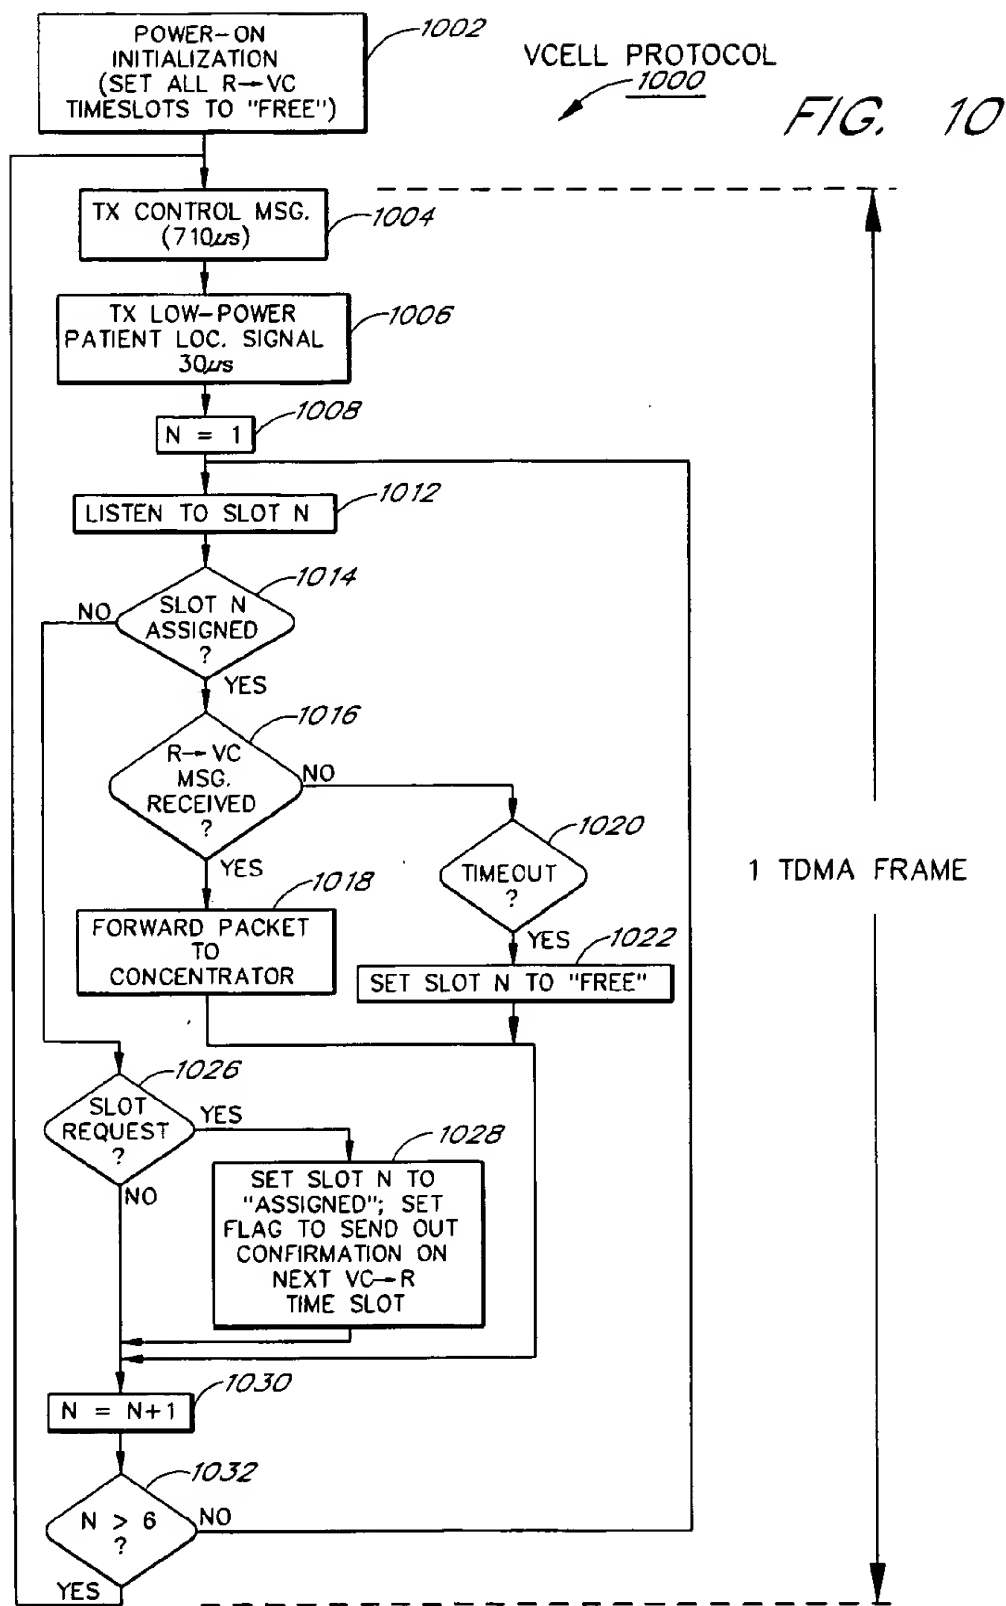
FIG. 9

VCELL CATALOG

1100

	RATING	CONNECTED TO	LOW-PWR. SIGNAL STRENGTH
VCELL1 (f1)	7		1
VCELL2 (f2)	8	✓	9
VCELL3 (f3)	8		6
VCELL4 (f4)	8	✓	6
VCELL5 (f5)	6		1
VCELL6 (f6)	2		0
VCELL7 (f7)	0		0
VCELL8 (f8)	0		0
VCELL9 (f9)	3		0
VCELL10 (f10)	5		1

FIG. 11



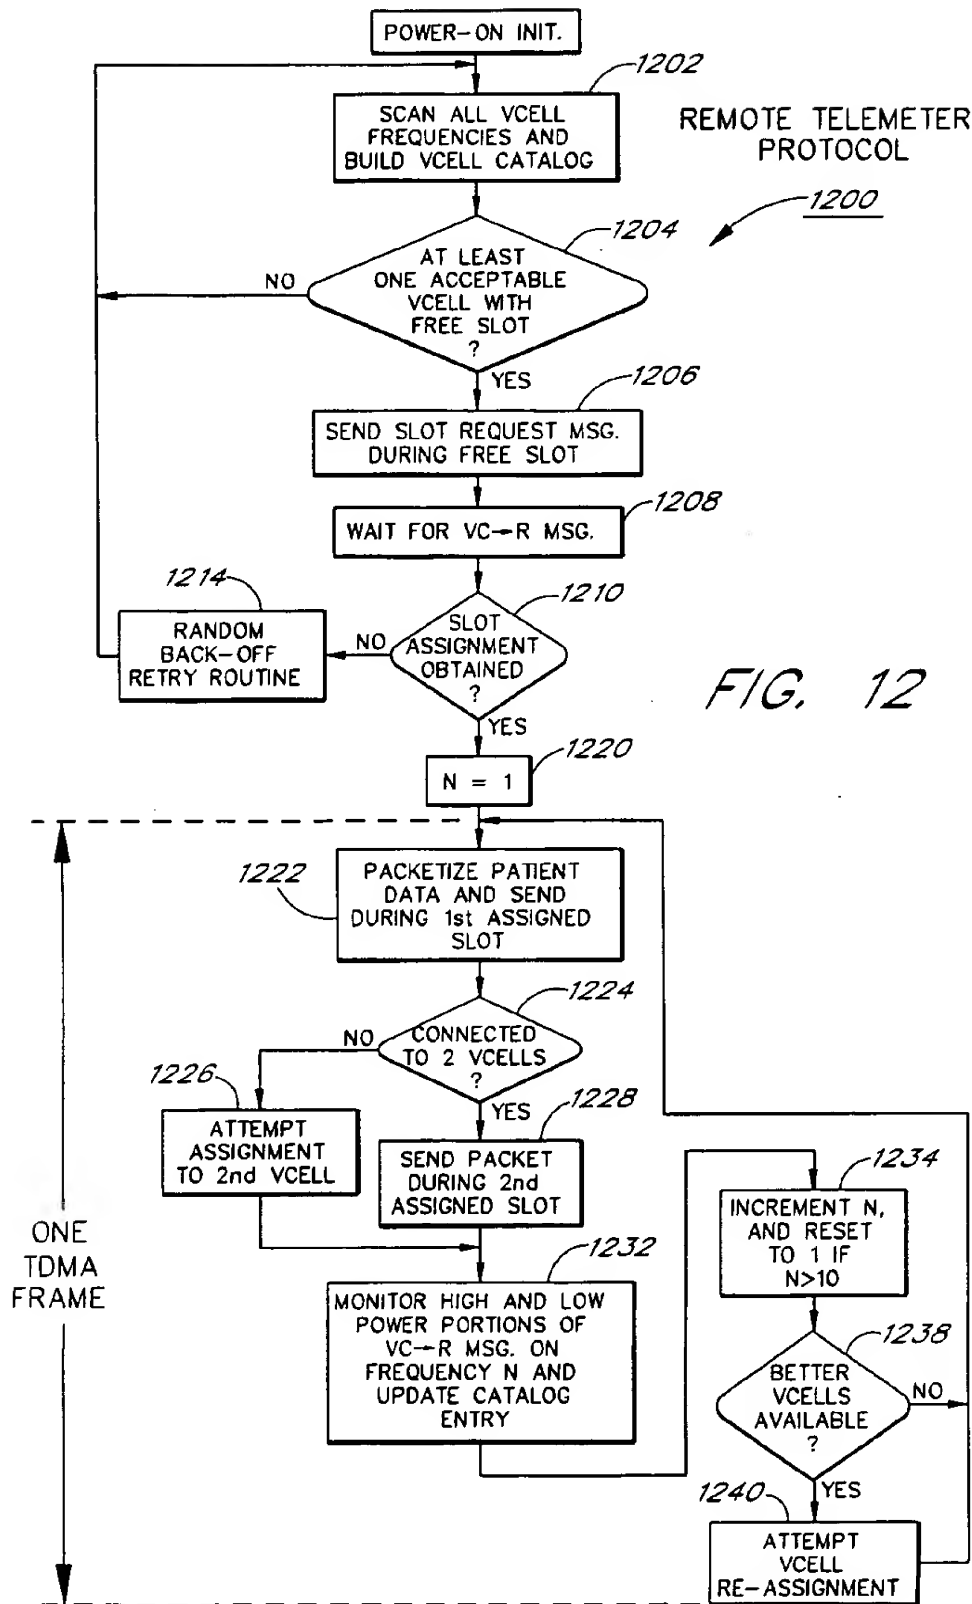
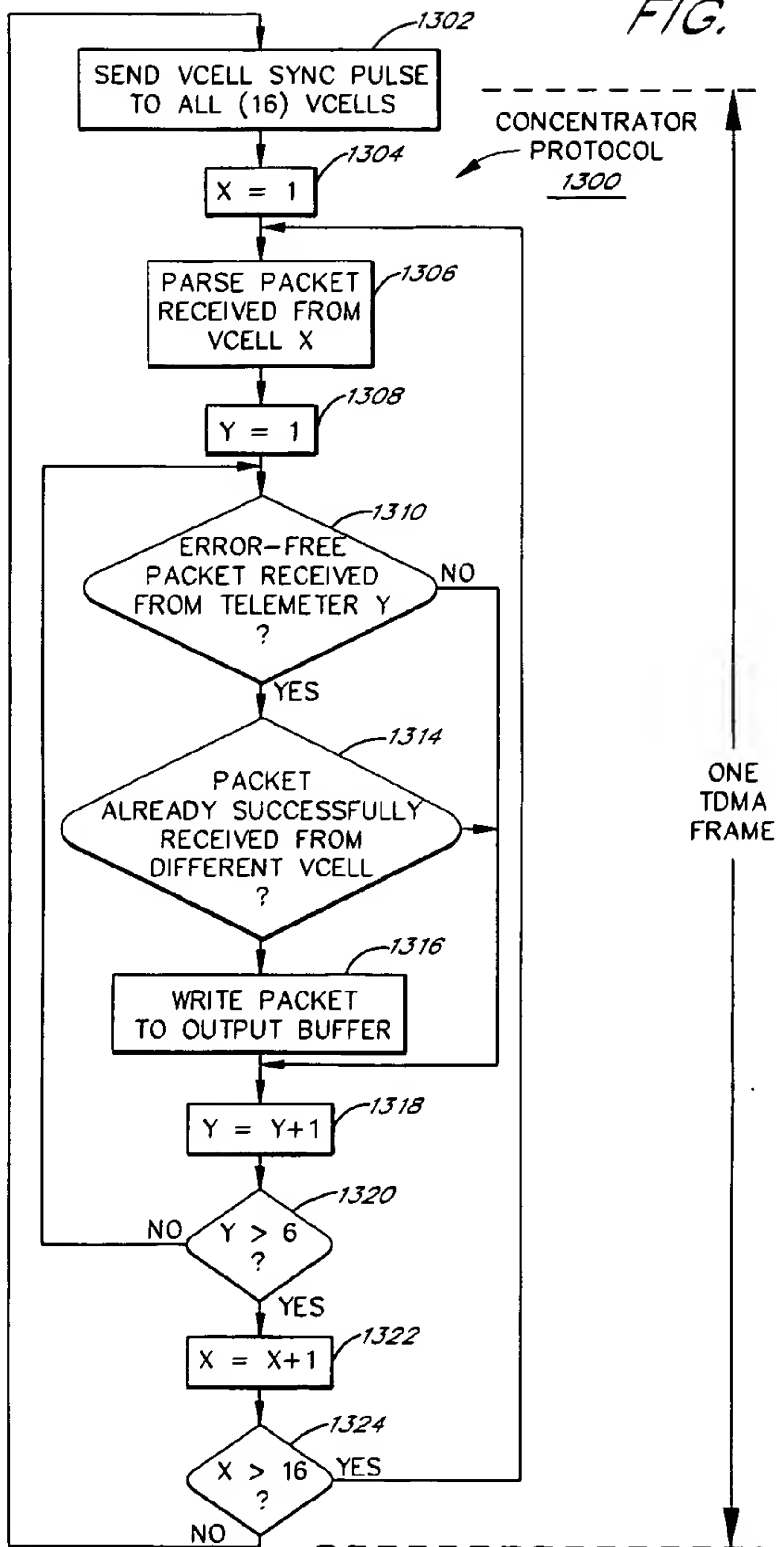


FIG. 13



TELEMETER DESIGN AND DATA TRANSFER METHODS FOR MEDICAL TELEMETRY SYSTEM

This application is a division of U.S. Appl. Ser. No. 08/675,594 filed Jul. 3, 1996, (now U.S. Pat. No. 5,944,659) which claims the benefit of U.S. Provisional Appl. Ser. No. 60/006,600 titled TWO-WAY TDMA TELEMETRY SYSTEM, filed Nov. 13, 1995.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to digital wireless communications systems of the type which employ portable, battery-powered communications devices, such as remote telemeters worn by ambulatory hospital patients for monitoring purposes. More particularly, the present invention relates to a network architecture, and an associated TDMA (time division multiple access) communications protocol, for facilitating the efficient and reliable exchange of information between portable wireless devices and centralized monitoring stations.

2. Description of the Related Art

Medical telemetry systems that allow the physiologic data of multiple, remotely-located patients to be monitored from a central location are known in the art. These systems typically comprise remote telemeters that remotely collect the physiologic data of respective patients and transmit the data over a wireless link to a centralized monitoring station. This physiologic data may include, for example, real-time electrocardiograph (ECG) waveforms, CO₂ levels, and temperature readings. From the centralized monitoring station, a clinician can visually monitor the physiologic status, in real time, of many different patients. The central station may also run automated monitoring software for alerting the clinician whenever a predetermined physiologic event occurs, such as a cardiac arrhythmia condition.

Remote telemeters of medical telemetry systems are generally of two types: instrument remote telemeters and ambulatory remote telemeters. An ambulatory remote telemeter is a portable, battery-powered device which permits the patient to be monitored while the patient is ambulatory. The ambulatory telemeter attaches to the patient by a strap or other attachment device, and receives the patient's physiologic data via ECG leads (and/or other types of sensor leads) which attach to the patient's body. The physiologic data is continuously transmitted to the central monitoring station by the telemeter's RF (radio frequency) transmitter to permit real-time monitoring. (A design of a remote transceiver which may be used in a two-way, ambulatory telemeter is described in the above-referenced provisional application.) Instrument remote telemeters operate in a similar manner, but receive the patient's physiologic data from a bedside monitor (or other instrument) over a hardwired link, such as an RS-232 connection. Instrument remote telemeters that transfer the physiologic data to the central station over a hardwired connection are also common.

SUMMARY

One problem that is commonly encountered in the field of medical telemetry involves signal loss caused by multi-path interference. Multi-path interference is a well-known phenomenon which occurs when a signal takes two or more paths (as the result of signal reflections) from the transmitter to the receiver such that the multi-path components destructively interfere with each other at the receiver's antenna. To

reduce the effects of multi-path interference, some telemetry equipment manufactures have included multiple antenna/receiver pairs on each remote telemeter. With this technique, known as spacial diversity, when one of the antennas experiences multi-path fading, the other antenna (and the corresponding receiver) is used to receive the signal. One problem with this method is that it adds to the cost, size and complexity of the remote telemeter. In addition, in at least some implementations, a loss of data may occur when a "switch-over" is performed from one antenna/receiver pair to the other.

Another problem that has been encountered in the field of medical telemetry relates to the ability to monitor a large number of patients over a coverage area that extends to all patient areas of the hospital. A common solution to this problem involves installing a large number of antennas (e.g., 200 or more) throughout the hospital (with different antennas positioned in different patient areas), and interconnecting the antennas using signal combiners to form a single, distributed antenna system. One problem with this "distributed antenna system" approach is that each antenna and its associated preamplifier (or preamplifiers) contributes to the noise floor of the antenna system, and thereby increases the minimum transmit power at which the transmitting components of the system can operate. (The reasons for this noise floor degradation are discussed below.) Consequently, unless the transmission power of the system's transmitters is increased, a practical limitation is imposed on the number of antennas that can be included in the system, and on the coverage area provided by the system.

Although the noise floor degradation problem can potentially be overcome by increasing the transmission power of the telemetry equipment, there are at least two problems associated with increasing the transmit power. The first problem is that under existing Federal Communications Commission (FCC) regulations, medical telemetry equipment is only permitted to operate within certain frequency bands, and must operate within certain prescribed power limits within these bands. Under FCC Part 15.241, for example, which governs the protected VHF (174-216 MHz) medical telemetry band (a band which is generally restricted to VHF television and medical telemetry), telemetry devices are not permitted to transmit at a signal level which exceeds 1500 microvolts/meter at 3 meters. To operate at power levels which exceed this maximum, frequency bands which offer less protection against interference must be used. The second problem is that increasing the transmit power of an ambulatory telemeter will normally produce a corresponding reduction in the telemeter's battery life.

Another problem with distributed antenna systems is that they are typically highly vulnerable to isolated sources of electromagnetic interference ("EMI"). Specifically, because the signals received by all of the antennas are combined using RF signal combiners, a single source of interference (such as a cellular phone or a faulty preamplifier) at or near one of the antennas can introduce an intolerable level of noise into the system, potentially preventing the monitoring of all patients. One consequence of this problem is that antennas generally cannot be positioned near known intermittent sources of EMI such as X-ray machines, CAT (computerized axial tomography) scanners, and fluoroscopy machines, preventing patient monitoring in corresponding diagnostic areas.

In light of these and other problems with existing medical telemetry systems, the present invention seeks to achieve a number of performance-related objectives. One such objective is to provide an architecture in which the coverage area

and patient capacity can be increased without degrading the noise floor. This would allow the telemetry system to be expanded in size and capacity without the need to increase the transmit power of the battery-powered remote telemeters, and without the need to operate outside the protected VHF medical telemetry band. A related objective is to provide an architecture which is highly scalable, so that the capacity and coverage area of the system can easily be expanded through time.

Another goal of the invention is to provide extensive protection against signal drop-outs caused by multi-path interference. The present invention seeks to achieve this objective without the need for multiple antennas or receivers on the telemeters, and without the loss or interruption of physiologic data commonly caused by antenna/receiver switch-overs. A related goal is to provide a high degree of protection against isolated sources of EM, and to allow patients to be remotely monitored while near known intermittent sources of interference.

Another goal of the invention is to provide an architecture in which a large number of patients (e.g., 500 to 800 or more) can be monitored using a relatively narrow range of RF frequencies (such as the equivalent of one or two VHF television channels). This would allow the RF communications components of the system to be optimized for narrow-band operation, which would in-turn provide a performance advantage over wide-band systems.

In accordance with these and other objectives, a medical telemetry system is provided which includes multiple remote telemeters (which may include both ambulatory and instrument telemeters) which transmit the real-time physiologic data of respective patients via RF to multiple ceiling-mounted transceivers, referred to as "VCELLs." The VCELLs are hardwire-connected to a real-time data distribution network which includes at least one centralized monitoring station. (In a preferred implementation, each group of 16 VCELLs is connected via twisted pair lines to a respective "concentrator PC," and the concentrator PCs and monitoring stations are interconnected as part of a hospital local area network.) The VCELLs are distributed throughout the hospital such that different VCELLs provide coverage for different patient areas, and are spaced such that the coverage zones provided by adjacent VCELLs overlap with one another. Different VCELLs within the same general area communicate with the remote telemeters on different respective RF frequencies (i.e., frequency channels), so that a remote telemeter can selectively communicate with a given VCELL by selecting that VCELL's frequency. As described below, however, VCELL frequencies are reused by VCELLs that are spaced sufficiently apart from one another to avoid interference, allowing the system to be implemented as a narrow-band system which uses a relatively small number of frequencies (e.g., 10) to provide coverage for an entire hospital facility.

In a preferred embodiment, the remote telemeters communicate with the VCELLs using a wireless time division multiple access (TDMA) protocol in which each VCELL can concurrently receive the real-time physiologic data of up to six remote telemeters (corresponding to six patients). As part of this protocol, the remote telemeters implement a VCELL "switch-over" protocol in which the telemeters establish wireless connections with different VCELLs based on periodic assessments (made by the telemeters) of the wireless links offered by the different VCELLs. Thus, as a patient moves throughout the hospital, the patient's remote telemeter may connect to (and disconnect from) many different VCELLs.

In operation, the remote telemeters send data packets (during assigned timeslots) to the respective VCELLs with which the telemeters have established wireless connections. (As described below, each remote telemeter preferably remains connected to two different VCELLs at-a-time to provide extensive protection against multi-path interference.) These data packets include the real-time physiologic data of respective patients, and include ID codes which identify the remote telemeters. The VCELLs in-turn forward the data packets to the real-time data distribution network to permit the real-time monitoring of the patients of the system.

To provide protection against multi-path interference and other causes of data loss, each remote telemeter maintains wireless connections with two different VCELLs at-a-time, and transmits each data packet to both of the VCELLs. These duplicate packet transmissions to the two different VCELLs take place on different frequencies during different TDMA timeslots. The two VCELLs forward the data packets to a centralized node (which may be a monitoring station or a concentrator PC in the preferred embodiment), which performs error correction by selecting between the corresponding packets based on error detection codes contained within the packets. Thus, the patient's physiologic data is sent from the remote telemeter to the centralized node over two separate data paths. Because the two VCELLs are spaced apart, and because the duplicate packets are transferred to the VCELLs on separate frequencies at different times, the packet transfers benefit from the protection offered by spacial diversity, frequency diversity and time diversity.

The architecture of the above-described medical telemetry system provides numerous advantages over prior art systems. One such advantage is that the system can be expanded in patient capacity and coverage area, by the addition of VCELLs, without increasing the noise floor of the system beyond the natural thermal noise floor. (This is because the data signals received by the VCELLs are multiplexed digitally at baseband, rather than being combined by RF analog signal combiners.) Thus, unlike distributed antenna telemetry systems, the noise floor does not impose an upper limit on the size of the system. Moreover, the architecture can accommodate a large number of patients (e.g., 500 to 800 or more) using a low maximum transmission power, such as the maximum transmit power permitted by the FCC for operation within the VHF medical telemetry band.

Another advantage is that the architecture is highly immune to isolated sources of EMI. A source of EMI (such as a cellular phone), for example, will typically contaminate the signals received by no more than one or two nearby VCELLs, as opposed to introducing noise into the entire system. (Because the remote telemeters connect to two VCELLs at-a-time, and automatically switch to different VCELLs when bad link conditions are detected, the contamination of one or two VCELLs will typically result in little or no loss of telemetry data.) One benefit of this immunity is that VCELLs can be installed within X-ray rooms and other radiological diagnostic rooms which contain intermittent sources of EMI, allowing patients to be monitored in such areas.

Another advantage of the architecture is that it permits the reuse of RF frequencies by VCELLs that are sufficiently spaced apart (by about 500 feet in a VHF implementation) to avoid interference with each other. By extending this concept, the present invention provides coverage for the entire facility using a relatively small number of frequencies

which fall within a relatively narrow frequency band. In a preferred VHF implementation, for example, it is estimated that a typical hospital can be covered using only 10 to 12 VCELL frequencies which fall within a frequency band that is equal in width to about two adjacent VHF television channels. This characteristic of the architecture advantageously allows the telemetry transceivers to be optimized (through the appropriate selection of transceiver components) for a relatively narrow band of frequencies, which in turn improves performance.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other features of the invention are described below with reference to the drawings of a preferred embodiment, which is intended to illustrate and not to limit the invention:

FIG. 1 is an architectural drawing of the hardware components of a medical telemetry system in accordance with the present invention.

FIG. 2 illustrates the attachment of an ambulatory remote telemeter to a patient of the system.

FIG. 3 illustrates the basic hardware components of the concentrator PCs and the ceiling-mounted transceivers (VCELLs) of FIG. 1.

FIG. 4 illustrates the basic hardware components of the ambulatory remote telemeters of FIG. 1.

FIG. 5A is a generalized circuit diagram of a transceiver which may be used in the VCELLs and remote telemeters of the telemetry system.

FIG. 5B illustrates the output of the phase-locked loop (PLL) chip of FIG. 5A during the locking of the transmit frequency of a remote telemeter.

FIG. 6 illustrates an increase in dynamic range achieved by the present system over prior art telemetry systems.

FIG. 7 illustrates how VCELLs operating on different frequencies may be arranged within a hospital hallway in accordance with the invention.

FIG. 8 illustrates a TDMA frame of a wireless TDMA protocol used for the transfer of information between the remote telemeters and the VCELLs of the system.

FIG. 9 illustrates the basic timeslot status information stored by each VCELL as part of the wireless TDMA protocol.

FIG. 10 is a flowchart of a protocol followed by each VCELL as part of the wireless TDMA protocol.

FIG. 11 illustrates the basic VCELL status information stored by each remote telemeter as part of the wireless TDMA protocol.

FIG. 12 is a flowchart of a protocol followed by each remote telemeter as part of the wireless TDMA protocol.

FIG. 13 is a flow chart of a protocol followed by the concentrator PCs for processing packets received from the VCELLs.

In the drawings, the left-most digit (or digits) of each reference number indicates the figure in which the item first appears. For example, an element with the reference number 310 first appears in FIG. 3, and an element with reference number 1100 first appears in FIG. 11.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

To facilitate a complete understanding of the invention, the description of the preferred embodiment is arranged within the following sections and subsections:

1. OVERVIEW

(i) GENERAL OPERATION

(ii) HARDWARE COMPONENTS

(iii) PATIENT CAPACITY AND DATA THROUGH-PUT

(iv) NOISE FLOOR IMPROVEMENT

(v) PROTECTION AGAINST ISOLATED EMI SOURCES

(vi) VCELL SPACING AND FREQUENCY REUSE

2. COMMUNICATIONS BETWEEN REMOTE TELEMETERS AND VCELLS

(i) OVERALL WIRELESS TDMA PROTOCOL

(ii) VCELL PROTOCOL

(iii) REMOTE TELEMETER PROTOCOL

3. COMMUNICATIONS BETWEEN VCELLS AND CONCENTRATORS

(i) PROCESSING OF TELEMETER COMMANDS

4. DATA TRANSFERS OVER LAN

5. VCELL LOAD MONITORING

6. TRANSCEIVER CIRCUIT AND OPERATION

7. CONCLUSION

1. Overview (FIGS. 1-7)

FIG. 1 illustrates the general architecture of a two-way medical telemetry system which operates in accordance with the present invention. The system, referred to herein as the "VCELL system," includes a number of wireless remote telemeters 102A, 102B which collect, packetize and transmit the physiologic data of respective hospital patients. (As used herein, the term "wireless" means that data is transferred to and/or from the device over a wireless medium.) The remote telemeters 102 may include both patient-worn (ambulatory) remote telemeters 102A which connect directly to the patient (as generally illustrated in FIG. 2), and instrument remote telemeters 102B which connect to a bedside or other patient monitor 104. The physiologic data transmitted by the remote telemeters may include, for example, real-time ECG signals, blood pressure readings, CO₂ levels, and temperature readings. The remote telemeters 102 may additionally sense and transmit various types of non-physiologic data, such as battery-level status data, ECG loose-lead status data, and patient location data. (The term "patient data" is used herein to refer collectively to the physiologic and non-physiologic data captured by the remote telemeters 102.)

The remote telemeters 102 communicate bi-directionally with a number of ceiling-mounted radio transceivers 106, referred to as "VCELLs," using a time division multiple access (TDMA) protocol. In one mode of operation, each VCELL 106 can communicate with up to six remote telemeters 102 at-a-time at a rate of 10 kilobaud (Kbaud) per telemeter. The VCELLs 106 are spaced apart from one another (typically by about 50 to 75 feet, depending upon expected patient density) throughout the hospital to provide a "cell-like" coverage area which consists of overlapping zones of coverage.

Different VCELLs 106 of the system operate (i.e., transmit and receive data) on a different RF frequency channels ("frequencies") within the VHF medical telemetry band (174-216 MHz). However, VCELLs that are sufficiently spaced apart to avoid interference with one another may operate on like frequencies, as described below. The VCELLs 106 and telemeters 102 of the preferred embodiment operate in compliance with the spectrum utilization and transmission power limitations of FCC Part 15.241. Although the system preferably operates within the VHF medical telemetry band, other suitable frequency bands may be used. In addition, although the system uses frequency

division multiplexing to separate the data transmissions to and from different VCELLs 106, other channel separation techniques can be used.

Although the remote telemeters 102 and VCELLs 106 shown in FIG. 1 are of the type which communicate by radio frequency (RF), the system may also include "hardwired" remote telemeters and VCELLs which communicate over hardwire connections. For purposes of this description, however, it may be assumed that the terms "remote teleme- 5 ter" and "VCELL" refer to RF devices, except where indicated otherwise.

With further reference to FIG. 1, the VCELLs 106 are connected by conventional shielded twisted pair lines 110 to concentrator PCs 112 ("concentrators"). In the preferred embodiment, each concentrator 112 can accommodate up to sixteen VCELLs 106. In a typical hospital installation, one concentrator 112 will service a single floor of the hospital. The concentrators 112 provide connectivity between the VCELLs 106 and a hospital local area network (LAN) 116. The LAN 116 serves as a real-time data distribution system for distributing the physiologic data of the patients with a known latency. The LAN 116 includes a 100 Mbit/second backbone 118 which is based on the 100BaseTx (Ethernet) protocol. (The term "backbone" refers generally to the transmission medium and the networking cards of the LAN.) 20 Alternative LAN protocols which could be used include ATM (Asynchronous Transfer Mode) and FDDI (Fiber Distributed Data Interface) and others.

The LAN 116 includes multiple monitoring stations 120 for allowing hospital personnel to remotely view and otherwise monitor the real-time physiologic data of the patients of the system. Each monitoring station 120 is preferably in the form of a standard 486 or Pentium based PC (personal computer) which runs conventional patient monitoring software, such as the VCOM (MPC 1100) patient monitoring software package available from VitalCom Incorporated. The patient monitoring software can also be loaded onto the concentrator PCs 112 so that the concentrators double as monitoring stations. The LAN 116 may also include one or more gateway computers 124 for connecting the LAN 116 to other networks, such as the Internet, to permit the exchange of patient information with other medical facilities and patient sites.

As will be apparent, the architecture illustrated in FIG. 1 provides for a high degree of scalability. The system can initially be installed, for example, as a single concentrator PC 112 which serves as the sole monitoring station for a set of 16 (or fewer) VCELLs, which may include both RF and hardwired VCELLs. With the addition of a LAN, new VCELLs 106 and concentrators 112 can be added to increase the patient capacity and/or coverage area of the system. (As described below, the architecture allows new VCELLs to be added to the system without a corresponding degradation in performance caused by noise.) Monitoring stations 120 can be added to the LAN 116 as needed to permit the remote viewing and monitoring of patient data from various locations within the hospital.

(i) General Operation

In operation, the remote telemeters 102 send data packets to individual VCELLs 106 using a wireless TDMA protocol. These packets include the patient data collected by the remote telemeters 102 (or by patient monitors connected to the remote telemeters), along with the ID codes of the respective telemeters 102. The VCELLs 106 forward these data packets to the corresponding concentrators 112, which in-turn broadcast the patient data on the LAN 116 (in real 60

time) for viewing and automated monitoring by the monitoring stations 120. The wireless TDMA protocol includes control timeslots for allowing the VCELLs to pass control information (e.g., synchronization information, commands, and timeslot assignments) to the remote telemeters 102. In addition, the protocol supports a patient location method (described below) for monitoring the remote location of each patient.

To support patient mobility, the VCELLs 106 and remote telemeters 102 implement a "switch-over" protocol in which the telemeters 102 continuously attempt to establish connections with those VCELLs which offer the best link performance. As part of this protocol, each remote teleme- 5 ter continuously assesses the quality of the RF link to each VCELL that is within range. The telemeters store this link assessment information within respective VCELL "catalogs" (described below), and periodically evaluate these catalogs to determine whether a switch-over to a new VCELL is desirable. When a remote teleme- 10 ter 102 determines that a VCELL is available (i.e., has an open timeslot) which offers better link performance than a current VCELL (i.e., a VCELL to which the teleme- 15 ter is currently connected), the remote teleme- 20 ter attempts to connect to the new VCELL. (As described below, this involves sending a timeslot request message to the selected VCELL 106, and then waiting for confirmation message from the VCELL.) If the connection is successfully established, the remote teleme- 25 ter 102 drops its connection to the current VCELL 106. Thus, a remote teleme- 30 ter 102 will normally connect to many different VCELLs 106 (including VCELLs of different concentrators 112) as the patient moves throughout the hospital. Transitions between VCELLs occur without interruption or loss of data, and are thus seamless from the viewpoint of the monitoring clinician.

To provide protection against dropouts caused by multi-path interference (and other types of interference), each remote teleme- 35 ter 102 attempts to maintain a connection with two VCELLs 106 at all times. (In other implementations, the remote telemeters 102 may connect to three or more VCELLs 106 to provide even greater protection against multi-path interference.) Whenever two VCELL connections are established, the remote teleme- 40 ter 102 transmits each of its data packets to both of the VCELLs. These redundant transfers take place on different frequencies during different TDMA timeslots. Thus, each wireless data path benefits from the protection offered by space, time, and frequency diversity. Upon receiving the redundant packets, the concentrator 112 to which the two VCELLs 106 are connected (assuming the VCELLs are connected to the same concentrator) uses error detection codes contained within the packets to discard bad packets, and to discard duplicate packets when both packets are successfully received.

In one implementation of the system, the remote telemeters 102 can only connect to the VCELLs 106 of one concentrator 112 at-a-time. In this implementation, each remote teleme- 45 ter 102 attempts to stay connected to the VCELLs of the current concentrator 112, and switches over to a different concentrator only when deemed necessary. In another implementation, the concentrators 112 of the system are maintained sufficiently synchronized with one another to allow each remote teleme- 50 ter 102 to connect to VCELLs of two different concentrators 112. When this situation occurs, the task of discarding duplicate packets automatically shifts to the monitoring stations 120.

The operation of the system is described in further detail in the following sections.

(ii) Hardware Components (FIGS. 3-5A)

FIG. 3 illustrates the basic components of the concentrators 112 and VCELLs 106 of the system. Each concentrator

112 comprises a generic PC having two RS-422 input-output (I/O) cards 302 and a 100BaseTx LAN card 304. The PC may, for example, be a Pentium-based PC with 16 megabytes of memory. (Additional memory and a display monitor will normally be provided if the concentrator 112 is to double as a monitoring station.) The RS-422 and 100BaseTx cards 302, 304 are standard AT size components which can be purchased off-the-shelf at computer stores.

Each RS-422 card 302 includes eight external (full duplex) I/Os which connect, respectively, to eight standard twisted pair lines 110. Each twisted pair line 110 connects to a respective VCELL 106. The twisted pair lines 110 are preferably shielded 140 Kbaud lines with RJ-45 connectors. As is conventional, each twisted pair line includes four wires: a transmit (TX) wire, a receive (RX) wire, a positive voltage (+) wire, and a negative voltage (-) wire. The (+) and (-) wires are used to provide power to the VCELLs 106, and the (RX) and (TX) wires are used for the transfer of data.

Each VCELL 106 is in the form of a microcontroller-based transceiver 308 coupled to an antenna 312. The specifications of a transceiver which may be used in the preferred embodiment are listed in Table 1. (A transceiver circuit which may be used within the VCELLs is illustrated in FIG. 5A, and is described below.) The transceiver 308 is coupled to random access memory (RAM) 310 for buffering packet data and storing various status information.

TABLE 1

VCELL TRANSCEIVER SPECIFICATIONS	
Operating Frequency	204-216 MHz
Frequency Tuning	100 KHz
Transmit Power	1500 μ V/meter @ 3 meters
Modulation Type	FSK
Modulation Rate	80 Kbaud
Deviation	\pm 50 KHz
Receive Sensitivity	-90 dBm (BER < .001)
Tx/Rx Switching Time	<10 μ s
Antenna	$\frac{1}{2}$ Wave Turnstile
Power Supply	6 to 12 VDC, < 100 ma

As depicted in FIG. 4, each remote telemeter 102A includes conventional sensor circuitry 402 for sensing and digitizing the patient data of a respective patient. (In instrument remote telemeters 102B of the type shown in FIG. 1, the sensor circuitry normally resides primarily within the patient monitor 104.) The sensor circuitry 402 is coupled to a microcontroller-based remote transceiver 404, which is in-turn coupled to a RAM 406 and an antenna 408. The sensor circuitry 402, remote transceiver 404 and RAM 406 are powered by one or more batteries 412. The specifications of a remote transceiver 404 which may be used in the preferred embodiment are listed in Table 2.

TABLE 2

REMOTE TELEMETER TRANSCEIVER SPECIFICATIONS	
Operating Frequency	204-216 MHz
Frequency Tuning	100 KHz
Transmit Power	1500 μ V/meter @ 3 meters
Modulation Type	FSK
Deviation	\pm 50 KHz
Modulation Rate	80 Kbaud
Receive Sensitivity	-90 dBm (BER < .001)
Tx/Rx Switching Time	<10 μ s
Antenna	$\frac{1}{2}$ Wave Turnstile
Power Supply	2 Alkaline Batteries, < 25 ma

Although the architecture of FIG. 1 is not tied to any particular transceiver implementation, the remote trans-

ceiver circuit disclosed in the above-referenced provisional application (diagram reproduced as FIG. 5A) is well-suited for use as both the VCELL transceiver 308 and the remote transceiver 404. With reference briefly to FIG. 5A, the circuit includes a microcontroller 502 (preferably a 17C42) coupled to an EEPROM 512 which includes a firmware program stored therein. In the remote telemeters 102, the firmware program implements the remote telemeter side of the wireless TDMA protocol (described below). Likewise, in the VCELLs 106, the firmware program implements the VCELL side of the wireless TDMA protocol (also described below). An overview of the transceiver circuit of FIG. 5A is provided below under the heading TRANSCEIVER CIRCUIT AND OPERATION.

(iii) Patient Capacity and Data Throughput

Each VCELL 106 can receive the patient data of six patients (i.e., six remote telemeters 102) at a sustained maximum data rate of 10 Kbaud per patient. (This data rate corresponds to one timeslot per TDMA frame using a simple FM transmitter, as described below.) In addition, the architecture supports increased data rates at the expense of reduced patient capacity. For example, a VCELL 106 could receive the patient data of three patients (i.e., three remote telemeters 102) at a data rate of 20 Kbaud per patient.

The total patient capacity of the system is limited primarily by the throughput of the LAN 116. In the preferred embodiment, the system design supports approximately 900 patients at a data rate of 10 Kbaud per patient 102. This results in a backbone throughput requirement of 900 \times 10 Kbaud=9 megabaud (Mbaud) at the network level. The use of 100BaseTx for the backbone supports this data rate while providing a margin of over 90% for overhead processing (such as synchronization, signaling, and background status keeping tasks). The patient capacity of the system can be increased by adding a second backbone 118 to the LAN 116 to provide dual 100 Mbaud data paths.

(iv) Noise Floor Improvement (FIG. 6)

One significant benefit of the VCELL architecture is that it overcomes the above-described noise-floor degradation problem encountered with distributed antenna systems, and thus allows the system to be expanded in capacity (through the addition of VCELLs) without a corresponding reduction in signal quality. To illustrate the noise-floor degradation problem encountered with distributed antenna systems, reference will be made to FIG. 6, which illustrates example dynamic range values for a single-antenna system (left-hand side) and a 400 antenna system with preamplifiers (right-hand side) operating within the VHF medical telemetry band.

In general, telemetry systems operate between two limits of transmitted signal strength: (i) the minimum signal that the receiver can detect above the thermal noise floor (which is the "natural" noise floor created by the normal movement of charged particles), and (ii) the maximum signal that the transmitter can provide at very close range (as experienced when the transmitter resides directly below the receiver's antenna.) As illustrated in FIG. 6, the minimum detectable signal level (based on the thermal noise floor) in the single-antenna telemetry system will typically be about -90 dBm. The maximum allowed signal level within the VHF medical telemetry band is 1500 microvolts/meter at 3 meters (as specified by FCC Part 15.241), which corresponds to a signal level of about -40 dBm with the transmitter located directly below a ceiling-mounted telemetry antenna. Thus, a

single-antenna medical telemetry system will have a dynamic range of about 50 dB.

As indicated above, the process of combining the RF signals of the antennas of a distributed antenna system has a loss associated with it. This loss results from the need for signal combiners, and from the large amount of coaxial cable required to interconnect the various antennas. To compensate for this loss, distributed antenna systems use preamplifiers, typically at the antenna sites, to boost the RF signal. One problem with this approach is that each preamplifier contributes to the noise level of the antenna system in excess of the noise actually received by the corresponding antenna. Thus, although only a few of the antennas typically receive a usable signal of a particular telemeter at any given time, all of the antennas (preamplifiers) contribute to the noise floor.

Consequently, each time the number of antennas of the distributed-antenna system is doubled, the minimum detectable signal level increases by about a factor of 2, or 3 dB. Thus, for example, a system with 400 antennas and 400 preamplifiers will suffer from a noise floor degradation of more than 24 dB (corresponding to over 8 doublings of the noise floor), producing a degraded dynamic range of less than 26 dB (FIG. 6). (A distributed antenna system of this size is currently installed at Barnes Hospital in St. Louis.) To reclaim this lost dynamic range, the remote telemeters could potentially be operated at a higher transmission power. However, the use of a higher transmission power would reduce the average battery life of remote telemeters. Moreover, an increase in power beyond the limits imposed by FCC Part 15.241 would require operation outside the protected medical telemetry band, exposing the system to new forms of RF interference.

In contrast to distributed antenna systems, the VCELL system combines the outputs of the VCELLs at baseband using digital multiplexing techniques. As a result, virtually no degradation of the noise floor occurs as VCELLs are added to the system, and the system enjoys the full 50 dB dynamic range regardless of the number of VCELLs. This has the effect of increasing the perceived transmitted power by about 24 dB (a 200 fold increase) over the 400 antenna system in the example above.

(v) Protection Against Isolated EMI Sources

Another benefit of the VCELL architecture is that it inherently offers a high degree of immunity against isolated sources of EMI (electromagnetic interference). In the above-described distributed antenna system, a single source of interference (such as an X-ray machine or a faulty copying machine) near one of the antennas can introduce an intolerable level of noise to the entire system, and prevent the monitoring of all patients of the system. In contrast, in the VCELL system, the interference source will only effect the operation of the VCELLs 106 that are sufficiently close to the source. Moreover, the contamination of one or two VCELLs by an isolated interference source will often have little or no impact on the ability to monitor patients in the area, since each remote telemeter 102 normally maintains data connections to two VCELLs 106, and automatically connects to a new VCELL source when a drop in the quality of a VCELL link is detected.

One benefit of this interference immunity is that it allows patients to be monitored near known intermittent sources of interference. For example, patients can be monitored within x-ray, fluoroscopy, and CAT-scan rooms by simply placing VCELLs in these areas.

(vi) VCELL Spacing and Frequency Reuse (FIG. 7)

The VCELLs are preferably mounted sufficiently close to one another such that each patient of the system will normally be within range of multiple VCELLs 106 at any given time. (For operation at maximum power within the VHF medical telemetry band, spacings of 50 to 75 feet are suitable.) Such an arrangement allows the remote telemeters 102 to maintain connections with two VCELLs at-a-time, as is desirable for mitigating the effects of multi-path interference.

One benefit of the architecture, however, is that it allows the VCELLs to be spaced as closely together as necessary to accommodate different patient densities. For example, a relatively large number of VCELLs (each operating on a different frequency) can be placed within a hospital cafeteria to accommodate the high patient densities which may occur during meal times. Because the remote telemeters 102 only attempt to connect to the VCELLs 106 that have open timeslots (as described below), the telemetry load during such high-density events is automatically distributed among the VCELLs.

Although it is possible to configure the system such that every VCELL 106 operates (i.e., transmits and receives data) on its own unique frequency, considerable performance benefits (described below) can be realized by re-using the same set of frequencies in different regions of the hospital. In general, two VCELLs can operate on the same frequency provided that they are sufficiently spaced apart to avoid interference with one another. For operation within the VHF medical telemetry band (at the maximum allowed signal strength), a separation of 500 feet between such VCELLs is more than adequate. By assigning like frequencies during the installation process to VCELLs that are spaced 500 feet (or greater) apart, it is estimated that 10 to 12 frequencies will be sufficient to provide coverage for a typical hospital.

FIG. 7 illustrates how a set of ten frequencies can be re-used in different sections of a hospital hall 700. As illustrated, a set of ten frequencies, f1-f10, can be used to cover a 500 foot section of the hall using ten corresponding VCELLs. (Each frequency symbol in FIG. 7 represents one VCELL.) The same ten frequencies can then be used to cover the next 500 foot section of the hallway. With appropriate staggering of frequencies between hospital floors, the same 10 frequencies can be used to provide coverage of an entire multi-floor hospital. (While ten frequencies may be adequate for many installations, the actual number of frequencies will depend upon such factors as the hospital floor plan, the telemeter transmission power, and the expected patient densities in the various patient areas.)

The ability to reuse frequencies provides several advantages over conventional frequency division multiplexed systems. One advantage is that a reduced number of clear frequencies need to be identified during the installation process. Another advantage is that the transmitters, receivers and antennas of the system can be optimized to operate over a much narrower band of frequencies. For a system which operates within the VHF medical telemetry band, for example, the VCELL frequencies can be selected to fall within a band of one or two VHF television channels (such as channels 12 and/or 13, which tend to have the lowest ambient noise), rather than spanning the entire 174 to 216 MHz range. It is estimated that such optimization will add a performance margin of 6 to 10 dB over existing telemetry equipment which operates over the entire 174 to 216 MHz range.

2. Communications Between Remote Telemeters and VCELLS (FIGS. 8-12)

(i) Overall Wireless TDMA Protocol (FIG. 8)

FIG. 8 illustrates a single frame of the wireless TDMA protocol used between the remote telemeters 102 and the VCELLs 106. The frame repeats every 5 milliseconds, and consists of seven timeslots: a 740 microsecond (μ s) VC \rightarrow R (VCELL to remote telemeter) timeslot and six 710 μ s R \rightarrow VC (remote telemeter to VCELL) timeslots. The VC \rightarrow R timeslots are used to broadcast information to the remote telemeters 102. (As described below, all VCELLs of the system are synchronized, and thus transmit at the same time.) The R \rightarrow VC timeslots are assigned by the VCELLs 106 to individual telemeters 102, and are used to transfer information from the assigned telemeters to the VCELLs. All timeslots terminate with a 10 μ s dead period, which is sufficient to allow the devices to switch between transmit and receive modes.

In the preferred embodiment, the remote telemeters 102 and VCELLs 106 transmit at a raw data rate of 80 Kbaud, which corresponds to a bit time of 12.5 μ s. At this data rate, a total of $(700 \mu\text{s/slot})/(12.5 \mu\text{s/bit})=56$ bits are transmitted during the data portion of each R \rightarrow VC timeslot. The first six of the 56 bit times are used for synchronization of the receiver, leaving 50 bits for the transfer of telemetry data (including error detection codes). Because this 50 bit message repeats every 5 milliseconds, or 200 times per second, the total telemeter-to-VCELL throughput for a single timeslot assignment is $200 \times 50 = 10,000$ bits/second, or 10 Kbaud. This throughput rate is obtained in the preferred embodiment using simple FM transceivers in the VCELLs and telemeters. As will be recognized by those skilled in the art, higher throughput rates can be achieved, at a greater expense, by using transceivers which use more sophisticated modulation techniques, such as BPSK and QPSK.

With further reference to FIG. 8, the VCELLs broadcast respective control messages to the remote telemeters 102 during the first 700 μ s of each VC \rightarrow R timeslot. (Because all VCELLs within range of one another transmit on different frequencies, each telemeter 102 can listen to the control message of only one VCELL at-a-time.) The control messages are used to transmit the following information to the telemeters:

Synchronization Sequences. The telemeters use these sequences to initially become synchronized and to maintain synchronization with the VCELLs 106.

VCELL-Specific Timeslot Assignment Status Data. For a given VCELL, this status data indicates which of the six R \rightarrow VC timeslots (if any) are unassigned, and thus available for use. The telemeters use this information to formulate timeslot request messages to the VCELLs.

Telemeter-Specific Timeslot Assignment Messages. A timeslot assignment message (or "confirmation" message) is transmitted in response to a timeslot request message from a specific telemeter, and serves as an acknowledgement to the telemeter that it has successfully acquired the requested timeslot.

Telemeter-Specific Commands. This is an optional feature which may be supported by certain remote telemeters 102. A command may be sent, for example, to instruct a telemeter to take a blood pressure reading, or to enter into special mode of operation.

VCELL ID Codes. Each VCELL transmits a unique ID code which is used for patient location.

During the last 30 μ s of each VC \rightarrow R timeslot, each VCELL transmits a low-power ($\frac{1}{4}$ -power in the preferred

embodiment), unmodulated signal to allow each remote telemeter 102 to estimate the location of the respective patient. (The use of a low-power, unmodulated signal for this purpose produces a more accurate VCELL-telemeter distance measurement.) Each remote telemeter 102 measures the signal strengths of the low-power transmissions of the various VCELLs (by listening to different VCELL frequencies during different TDMA frames), and maintains a table (discussed below) of the detected signal strengths. As a low duty cycle task (e.g., once every 5 seconds), each telemeter 102 evaluates its respective table to estimate the closest VCELL (i.e., the VCELL with the greatest signal strength). Whenever a change occurs in the closest VCELL, the telemeter transmits the ID of the new VCELL to the hospital LAN 116 (FIG. 1). The monitoring stations 120 use this information to keep track of the locations of the patients of the system. In other embodiments of the invention, patient location may be accomplished by having the VCELLs periodically attach VCELL identification codes to the data packets received from the remote telemeters 102.

With further reference to FIG. 8, the six R \rightarrow VC timeslots are used by the remote telemeters 102 to transmit data packets to individual VCELLs 106. (As described below, each telemeter 102 transmits to only one VCELL at-a-time.) These data packets are generally of two types: (i) telemetry data packets which include the patient data (including patient location data) of individual patients, and (ii) timeslot request messages for requesting timeslot assignments. Once a R \rightarrow VC timeslot has been assigned by a VCELL to a remote telemeter, the remote telemeter has exclusive use of the timeslot until the telemeter disconnects. Once the telemeter disconnects, the VCELL modifies its control message (transmitted on the VC \rightarrow R timeslot) to indicate that the timeslot is available for use.

During normal operation, each R \rightarrow VC timeslot of a given VCELL 106 will be assigned, if at all, to a different remote telemeter 102. Thus, when all six R \rightarrow VC timeslots of the VCELL are assigned, the VCELL receives the telemetry data of six different remote telemeters 102. In other modes of operation, multiple timeslots of a single VCELL can be assigned to the same telemeter to allow the telemeter to achieve a higher data throughput rate.

(ii) VCELL Protocol (FIGS. 9 and 10)

The VCELL side of the above-described wireless TDMA protocol is implemented via a firmware program which is executed by the microcontroller of each VCELL 106. With reference to FIG. 9, this program maintains a timeslot status table 900 in VCELL RAM 310 to keep track of the assignment status of each R \rightarrow VC timeslot. As illustrated in FIG. 9, the information stored within the table 900 includes, for each timeslot, whether or not the timeslot is currently assigned or unassigned. In addition, for the timeslots that are assigned, the table indicates the number of consecutive frames that have passed without receiving an error-free data packet from the assigned telemeter; each VCELL uses this information to implement a timeout procedure to determine whether the assigned telemeter 102 has disconnected.

FIG. 10 is a flow chart which illustrates the VCELL portion of the wireless protocol. With reference to block 1002, during a power-on initialization sequence, the VCELL 106 updates its timeslot status table 900 to set all of the R \rightarrow VC timeslots to the "free" (unassigned) state. The VCELL then enters into a primary program loop which corresponds to a single TDMA frame. Referring to blocks 1004 and 1006 of this loop, during the VC \rightarrow R timeslot the VCELL transmits the 710 μ s control message followed by

15

the 30 μ s patient location signal, as illustrated in FIG. 8. This control message includes the timeslot assignment data (for all six R \rightarrow VC slots) stored in the timeslot status table 900. On the first pass through this loop following power-on, the control message will indicate that all six R \rightarrow VC timeslots are available for use. With reference to block 1008, the VCELL 106 then sets a slot counter (N) to one (corresponding to the first R \rightarrow VC timeslot), and enters into a sub-loop (blocks 1012-1032) for processing the data packets transmitted by the telemeters.

During each R \rightarrow VC timeslot, the VCELL attempts to receive any telemetry data packet transmitted during the timeslot (block 1012), and checks the timeslot status table 900 to determine whether the slot is assigned (block 1014). With reference to blocks 1016-1022, if the timeslot is assigned and a data packet was successfully received, the VCELL forwards the packet to the concentrator 112, and clears the corresponding "missed packets" counter in the status table 900. (The data packet is actually sent to the concentrator 112 following the TDMA frame as part of a larger "VCELL packet" which represents the entire frame.) If, on the other hand, the timeslot is assigned but no packet was successfully received, the VCELL 106 updates the timeslot status table 900 to indicate that a packet was missed; in addition, the VCELL determines whether the number of consecutive missed packets has exceeded a timeout threshold (e.g., 64 packets). If the threshold is exceeded, the status table 900 is updated to set the timeslot to the "free" state.

With reference to blocks 1026 and 1028, if the timeslot is unassigned, the VCELL 106 determines whether it received a valid timeslot request message. If a valid timeslot request was received, the VCELL updates its status table to indicate that the slot has been assigned; in addition, the VCELL sets a flag to indicate that a timeslot confirmation message should be transmitted to the requesting telemetry 102 during the next VC \rightarrow R timeslot.

With reference to blocks 1030 and 1032, once all processing of the received telemetry packet (if any) is performed, the VCELL increments its timeslot counter. If the incremented counter is 6 or less, the program loops back to block 1012 to begin receiving any data transmitted during the next timeslot. If the incremented counter has exceeded 6, the program loops back to block 1004 to transmit the next control message.

Although the FIG. 10 flowchart illustrates the above-described operations in sequential order, it will be recognized that some of these operations can (and normally will) be performed concurrently or out-of order. For example, the step of forwarding the data packets to the concentrator (block 1018) is preferably performed as a separate task, with all of the telemetry packets received during the TDMA frame transferred together within a larger VCELL packet.

(iii) Remote Telemetry Protocol (FIGS. 11 and 12)

The telemetry side of the wireless TDMA protocol generally mirrors the VCELL protocol 1000, and is similarly implemented by a firmware program which is executed by the microcontroller of each remote telemetry 102. As part of this protocol, each remote telemetry 102 maintains a VCELL catalog within its respective RAM 406. The general format of this catalog is illustrated in FIG. 11, which is representative of a system which uses a total of 10 VCELL frequencies. As illustrated in FIG. 11, the catalog 1100 includes one set of entries for each of the ten VCELL frequencies. The telemetry thus stores status information for up to ten nearby

16

VCELLs at-a-time. The entries stored with respect to each VCELL frequency include the following:

Rating. A rating of the quality of the RF link to the VCELL 106, as assessed by the individual telemetry 102. The RF links are assessed by the telemeters one frequency at-a-time during the control message portions of the VC \rightarrow R timeslots. (In other embodiments, the task of assessing the available RF links may alternatively be performed by the VCELLs.) In one embodiment, the VCELL ratings are based on a combination of signal strength (as measured by the telemeters) and bit error rate. Specifically, if an error is detected in a VCELL's transmission, the VCELL's rating is set to a zero or null entry to indicate that the VCELL should not be used; and if no error is detected, the rating is set to a value which is proportional to the measured signal strength. The ratings are periodically compared (as described below) to determine whether to attempt a connection to a new VCELL.

Connected To. A flag indicating whether or not the telemetry is connected to a VCELL on the frequency. During normal operation, each telemetry will be connected to two different VCELLs (on two different frequencies) at-a-time.

Low-Power Signal Strength. A measurement of the signal strength taken during the low-power (patient location) portion of the VC \rightarrow R timeslot. The low-power signal strengths stored in the catalog 1100 are periodically compared to estimate which of the VCELLs the patient is closest to. When the outcome of this comparison changes, the telemetry transmits the new location (i.e., the VCELL ID, which is obtained from the VCELL during the high-power message portion) to the hospital LAN 116 in a subsequent data packet.

The remote telemeters 102 also keep track of the unique IDs of the VCELLs that are within range.

The protocol followed by the remote telemeters 102 generally consists of the following steps:

1. Scan all VCELL operating frequencies and construct the VCELL catalog 1100. Remain in this mode until at least one VCELL 106 is identified which has an acceptable rating and an unassigned timeslot.
2. Send a timeslot request message to the VCELL identified in step 1 during one of the available timeslots. Perform this task using a random back-off algorithm in case other remote telemeters attempt to connect to the VCELL during the same timeslot. Remain in this operating mode (including step 1) until a timeslot assignment message is received from the selected VCELL.
3. Once connected to a first VCELL ("VCELL 1"), attempt to connect to the "next best" VCELL ("VCELL 2") in the catalog which has an acceptable rating and a free timeslot (other than the timeslot being used to communicate with VCELL1), to provide a second (diversity) data path. Send telemetry data packets to VCELL1 during this process.
4. Monitor the catalog entries to determine whether any of the other VCELLs offer better link performance than VCELLs 1 and 2. When a better VCELL is available (i.e., has an open, nonconflicting timeslot), send a timeslot request message to the "new" VCELL. Drop the connection with the current VCELL once a timeslot assignment message is received.
5. As a background task, scan the VCELL frequencies and update the catalog. This can be done as a low priority, low duty cycle task (such as once per second).

17

FIG. 12 illustrates this protocol in greater detail for a system which uses 10 VCELL frequencies. With reference to block 1202, the remote telemeter 102 initially scans the ten VCELL frequencies and builds the VCELL catalog 1100. This involves monitoring different VCELL frequencies during different TDMA frames. With reference to blocks 1204-1210, once an acceptable VCELL 106 with an open timeslot has been identified, the telemeter 102 transmits a timeslot request message to the selected VCELL, and then monitors the selected VCELL's frequency during the following VC→R timeslot to determine whether the slot has been successfully acquired. With reference to block 1214, if the connection attempt is unsuccessful, a random back-off algorithm (such as the binary exponential back-off algorithm used by Ethernet) is used to retry the connection attempt. If the retry is unsuccessful (after, for example, 2 retry attempts), or if it is determined that the requested timeslot has been assigned to a different telemeter 102, the telemeter repeats the above process to identify another potential VCELL.

With reference to blocks 1220-1226, once a timeslot assignment has been obtained from a first VCELL, the protocol enters into a loop (blocks 1222-1240) which corresponds to a single TDMA frame. Within this loop, the telemeter 102 sends telemetry data packets to the first VCELL (block 1222) during the assigned timeslot while attempting to connect to a second VCELL (block 1226). The process of connecting to the second VCELL is generally the same as the above-described process for connecting to the first VCELL, with the exception that the timeslot used to communicate with the second VCELL must be different from the timeslot used to communicate with the first VCELL. With reference to block 1228, once a second VCELL connection has been established, the telemeter sends all data packets to both VCELLs.

With reference to blocks 1232 and 1234, the remote telemeter 102 monitors the high-power and low-power transmissions of the VCELLS 106 (during the VC→R timeslots) and updates the rating and low-power signal strength entries of the VCELL catalog 1100. Although this process is shown in FIG. 12 as occurring during every TDMA frame (one VCELL per frame), this function can alternatively be performed as a low duty cycle task.

With reference to blocks 1238 and 1240, as a background task the telemeter 102 monitors the VCELL catalog 1100 to determine whether a VCELL 106 with a higher rating exists. If a VCELL with a higher rating and an available (nonconflicting) timeslot is identified, the telemeter attempts to connect to the new VCELL (as described above), and if successful, drops the existing connection to one of the two "current" VCELLs.

As a background, low priority task, the telemeter 102 also monitors the VCELL catalog 1100 to determine whether a change has occurred in the VCELL 106 with the greatest low-power signal strength. (This process is omitted from FIG. 12 to simplify the drawing.) When such a change occurs, the telemeter transmits the VCELL ID of the new "closest" VCELL in a subsequent telemetry packet. As described above, this information is used by the monitoring stations 120 to track the location of each patient.

3. Communications Between VCELLS and Concentrators (FIG. 13)

The VCELLS 106 and concentrators 112 communicate bi-directionally over the shielded twisted pair lines 110 in accordance with the RS-422 specification. (RS-422 is an Electronic Industries Association interface standard which defines the physical, electronic and functional characteristics

18

of an interface line which connects a computer to communications equipment.) The RS-422 interface supports an overall data transfer rate of 140 Kbaud, which corresponds to 20 Kbaud per TDMA timeslot.

In operation, each VCELL 106 sends one packet to its respective concentrator 112 for every TDMA frame; this VCELL packet includes all of the telemeter packets (up to six) received during the corresponding TDMA frame. (The terms "VCELL packet" and "telemeter packet" are used in this description to distinguish between the two types of packets based on their respective sources.) The concentrator 112 in-turn parses the VCELL packet to extract the individual telemeter packets, and performs error checking on the telemeter packets using the error detection codes contained within such packets. As part of the error checking protocol, the concentrator 112 discards all telemeter packets which include errors (or uncorrectable errors if error correction codes are used), and discards all telemeter packets that are redundant of packets already received from a different VCELL. All other packets are written to an output buffer for subsequent broadcasting over the LAN 116.

FIG. 13 is a flow chart which illustrates the concentrator side of the VCELL-to-concentrator protocol in further detail. With reference to block 1302, the concentrator 112 sends a VCELL synchronization pulse to its 16 VCELLs once per TDMA frame. The concentrator then initializes a loop counter (block 1304), and enters into a loop (blocks 1306-1324) in which the concentrator processes the 16 VCELL packets (one per loop) received from the 16 VCELLs. Within this loop, the concentrator 112 parses each VCELL packet (block 1306) to extract the individual telemeter packets contained therein, and then enters into a sub-loop (blocks 1310-1320) in which the concentrator performs error checking (as described above) on the individual telemeter packets (one telemeter packet per sub-loop).

With reference to blocks 1310-1316, error free telemeter packets which have not already been successfully received (from other VCELLs) are written to an output buffer of the concentrator 112. (As described below, a separate concentrator task reads these packets from the buffer and broadcasts the packets on the LAN 116 during patient-specific timeslots of the LAN protocol.) With reference to blocks 1320-1324, once all of the telemeter packets within a given VCELL packet have been processed, the protocol loops back to block 1306 (unless all 16 VCELL packets have been processed, in which case a new synchronization pulse is transmitted), and the concentrator 112 begins to process the next VCELL packet.

As indicated by the foregoing, the concentrators 112 only place error-free telemeter packets on the LAN backbone 118, and do not place duplicate error-free telemeter packets (from the same telemeter) on the backbone 118. Nevertheless, the duplicate (error-free) packets transmitted by a telemeter will normally appear on the LAN backbone 118 when the remote telemeter connects to VCELLs of two different concentrators 112 (as permitted in one implementation of the invention). In this situation, the monitoring stations 120 simply ignore the extra telemeter packets.

(i) Processing of Telemeter Commands

In system implementations which support the sending of commands to the remote telemeters 102, the concentrators 112 additionally implement a simple task (not illustrated in FIG. 13) for receiving commands from the monitoring stations 120 and forwarding these commands to the VCELLS 106. As part of this task, each concentrator 112

maintains a list of all of the remote telemeters 102 to which the concentrator is currently connected. (This list is generated by monitoring the telemeter ID codes contained within the telemeter data packets.) When a monitoring station 120 places a telemeter-addressed command on the LAN 116, each concentrator 112 of the system receives the command and checks its respective list to determine whether a connection exists with the target telemeter. If a concentrator determines that such a connection currently exists, the concentrator 112 sends the command to all 16 of its VCELLS 106. The sixteen VCELLs in-turn transmit the telemeter command during a subsequent VC→R timeslot. To increase the probability of receipt, the VCELLs are preferably configured to re-transmit the telemeter command over several TDMA frames. In other embodiments, an acknowledgement protocol can be implemented in which the telemeters embed an acknowledgement message within a subsequent data packet.

4. Data Transfers Over LAN

Data transfers over the LAN backbone 118 are accomplished using a real-time TDM (time division multiplexing) protocol which makes use of the 100BaseTx protocol. Among other things, this protocol distributes the telemetry data from the concentrators 112 to the monitoring stations 120 with a known latency, permitting the real-time monitoring of patient data.

Each 50 millisecond frame of the TDM protocol includes 1000, 50 μ s timeslots. Every remote telemeter 102, VCELL 106, concentrator 112, monitoring station 120, and gateway 124 of the system is uniquely assigned one of the 1000 backbone timeslots. The backbone timeslots that are uniquely assigned to respective remote telemeters 102 are used to transfer telemetry data packets (containing patient-specific physiologic data) from the concentrators 112 to the monitoring stations 112. All other entities that are connected to the LAN backbone 118 also have access to this telemetry data. The remaining backbone timeslots are used for the transfer of synchronization and control information between the various LAN entities.

The 100BaseTx backbone 118 has the capacity to transfer up to 5000 bits in each 50 μ s backbone timeslot. Thus, each remote telemeter (and other entity which is assigned a backbone timeslot) is effectively allocated a LAN bandwidth of 5000 bits/slot \times 20 frames/second = 100 Kbaud. This more than satisfies the 20 Kbaud data rate per telemeter which is required when a telemeter connects to VCELLs of two different concentrators.

Upon initialization of the system, a "master" concentrator 112 transmits a synchronization packet to all other concentrators of the LAN 116. This synchronization packet defines the starting point of the backbone TDM frame. Thereafter, the frame repeats at a rate of 20 frames per second. As indicated above, a task which runs on each concentrator moves telemeter packets from the concentrator's output buffer to the LAN during the appropriate patient-specific (50 μ s) timeslots. This task waits for a patient (telemeter) timeslot, and then transmits all corresponding telemeter packets which have been written to the output buffer since the same timeslot of the immediately preceding backbone frame. When a telemeter is connected to VCELLs of two different concentrators, the 50 μ s patient timeslot is divided equally between the two concentrators. This is accomplished by passing control messages between the concentrators.

5. VCELL Load Monitoring

As a background task, each concentrator 112 maintains a statistical log or "histogram" of the loads carried by each of

the concentrator's VCELLs 106. This histogram can periodically be examined by network administrators to evaluate the current positioning of the VCELLs. When, for example, the histogram indicates that a VCELL in a particular patient area reaches its capacity (i.e., all six timeslots assigned) on a frequent basis, another VCELL (which operates on a different frequency) can be installed in the area to reduce the load on the heavily-loaded VCELL.

6. Transceiver Circuit and Operation (FIGS. 5A and 5B)

The transceiver circuit illustrated in FIG. 5A will now be described. As indicated above, this general circuit can be used in both the remote telemeters 102 and the VCELLs 106 of the system. When included within a remote telemeter 102, the transceiver will typically be powered by battery. When included within a VCELL 106, the transceiver will be powered by the corresponding concentrator 112 over a twisted pair line 110 (as illustrated in FIG. 3).

The transceiver 112 comprises a microcontroller (preferably a 17C42) which is connected, via appropriate port lines, to a programmable phase-locked loop chip 504 ("PLL chip"), a voltage controlled oscillator (VCO) 506, a receiver (RCVR) 508, a set of DIP (dual in-line package) switches 510, an EEPROM 512, and a sample-and-hold (S/H) device 520. (The sample-and-hold 520 is preferably omitted in the VCELL transceivers 308.) The PLL chip 504 is preferably a Motorola MC 145192 which can be placed, via appropriate commands, into a low-power state when not in use. The microcontroller 502 is clocked by an 8 MHz high stability ($\pm 0.001\%$) crystal oscillator 516. The output of the amplifier 524 and the signal input of the receiver 508 are connected to respective terminals of a transmit/receive switch 528, which is connected to the antenna 312, 408 via a band-pass filter (BPF) 530.

As illustrated in FIG. 5A, the PLL chip 504 is coupled to the VCO 506 to form a phase lock loop circuit. Via the PLL chip 504, the phase lock loop circuit can be programmed to generate a carrier signal of a selected frequency. Within the transceivers of the telemeters 102, the sample-and-hold device 520 is connected so as to allow the microcontroller 502 to programmably interrupt the phase-lock process and hold the carrier frequency at a steady frequency value within a preselected margin of frequency error. This allows the carrier frequency to be locked rapidly, at low power, without waiting for a phased-locked state to be reached. In a preferred embodiment, this feature is used to lock the transmit frequency of each telemeter just prior to each R→VC timeslot to which the telemeter is assigned.

As illustrated in FIG. 5A, the VCO 506, amplifier 524, and receiver 508 are coupled to the power supply (V_{SUP}) via respective microcontroller-controlled switches 534, 536, 538 such that the microcontroller 502 can selectively turn these components ON and OFF to conserve power. (In the VCELLs, the switches 534, 536, 538 can be omitted since battery life is not a concern.) To utilize this power-conservation feature, the firmware program (stored within the EEPROM 512) of the remote telemeters 102 includes code for maintaining these active transceiver components in an OFF state when not in use. For example, the receiver is maintained in an OFF state during TDMA timeslots for which the remote telemeter 102 is not receiving data, and the amplifier is maintained in an OFF state during timeslots for which the remote telemeter 102 is not transmitting. This feature of the transceiver circuit significantly increases the average battery life of the remote telemeters 102.

In operation within a remote telemeter, the microcontroller 502 maintains the PLL chip 504 in its low-power state,

21

and maintains the amplifier 424, VCO 506 and receiver 508 in respective OFF states, during timeslots for which the telemeter is neither transmitting nor receiving data. Shortly before the next R→VC timeslot which is assigned to the telemeter 102, the microcontroller 502 initiates a frequency lock operation which involves initiating a phase-lock process, and then interrupting the process (by opening the sample-and-hold 520) once the carrier frequency has settled to within an acceptable margin of error. This process is illustrated in FIG. 5B, which is an approximate graph of the output (V_{PLL}) of the PLL chip 504 following power-up at T_0 .

With reference to FIG. 5B, just prior to T_0 , the VCO 506 is turned on, the sample-and-hold 520 is in the closed (or "sample") position, and the PLL chip 504 is in the low-power state. At T_0 , the PLL chip 504 is taken out of the low-power state, causing its output V_{PLL} to ring, and thus causing the output of the VCO to oscillate above and below the programmed transmit frequency. Following T_0 , the output of the PLL is in the general form of a damped sinusoid, which approaches the voltage that corresponds to the programmed frequency. (Because the voltage V_{PLL} controls the VCO 506, the amplitude of the voltage signal in FIG. 5B corresponds to the frequency.)

Once this oscillation is sufficiently attenuated such that the frequency error is within a predetermined tolerance (e.g., ± 5 KHz), the sample-and-hold 520 is opened (at T_1 in FIG. 5) to hold the input voltage to the VCO 506. (This is accomplished by waiting a predetermined delay, T_{DELAY} , before opening the sample-and hold 520, as described in the above-referenced priority application.) This holds the output frequency, and ensures that the remote telemeter's subsequent data transmission will not be contaminated by any oscillation in the PLL's output. Immediately following T_1 , the amplifier 524 is turned on, the PLL 504 is placed in the low-power state, and the T/R switch 528 is placed in the transmit position. The microcontroller 402 then begins sending its transmit data to the VCO, to thereby FSK-modulate the carrier signal. Following the transmission of the telemeter packet, the amplifier 524, and VCO 506 are turned off.

In one embodiment, the telemeter firmware is written such that the telemeters 102 only use non-adjacent (i.e., non-consecutive) R→VC timeslots. This ensures that each telemeter will have at least a 720 μ s "dead period" between transmissions during which to lock the new transmit frequency using the above-described process.

The process of receiving data (during VC→R timeslots) is generally analogous to the above-described transmit process, with the exception that the sample-and-hold 520 is left in the closed (sample) position throughout the VC→R timeslot.

While the present invention has been described herein with reference to a preferred embodiment of a medical telemetry system, the invention is not so limited, and should be defined only in accordance with the following claims.

What is claimed is:

1. A remote telemeter for use in a medical telemetry system which supports real-time monitoring of ambulatory patients, comprising:

a processor which receives and processes real-time physiologic data of a patient, the physiologic data measured by sensors which attach to the patient;

a battery-powered transceiver responsive to the processor to transmit the physiologic data in data packets to selected receivers of a plurality of receivers, the plurality of receivers spatially distributed throughout a medical facility and coupled to a monitoring system

22

such that different receivers provide data reception coverage for different areas of the medical facility, the transceiver switchable by the processor between a plurality of wireless channels that correspond to the plurality of receivers; and

a control program executed by the processor to implement a wireless communications protocol in which the transceiver transmits the physiologic data to multiple different receivers of plurality on different, respective wireless channels and during different timeslots to provide redundant transmission paths, the redundant transmission paths providing spacial diversity, frequency diversity and time diversity.

2. The remote telemeter according to claim 1, wherein at least some of the receivers are transceivers that communicate bi-directionally with the remote telemeter.

3. The remote telemeter according to claim 1, wherein the plurality of wireless channels fall within the VHF medical telemetry band.

4. The remote telemeter according to claim 1, wherein the control program implements a switch-over protocol in which the remote telemeter establishes wireless links to selected receivers of the plurality of receivers in response to movement of the patient throughout the medical facility.

5. The remote telemeter according to claim 4, wherein the switch-over protocol switches between the receivers based at least upon assessments of wireless links to individual receivers of the plurality.

6. The remote telemeter according to claim 1, wherein the control program monitors patient location signals transmitted by the receivers, and uses the patient location signals to estimate a current location of the patient within the medical facility.

7. The remote telemeter according to claim 1, wherein the processor and transceiver are packaged within a housing which attaches to an ambulatory patient.

8. A remote telemeter for use in a medical telemetry system which supports real-time monitoring of patients that are mobile within a medical facility, comprising:

a transceiver circuit which communicates bi-directionally over a wireless channel with a plurality of transceivers that are coupled to a monitoring station, the plurality of transceivers spatially distributed throughout the medical facility such that different transceivers provide data reception coverage for different areas of the medical facility, the transceiver circuit configured to receive real-time physiologic data collected from a patient and to transmit the physiologic data to selected transceivers; and

a processor which implements a transceiver switch-over protocol to cause the transceiver circuit to switch between transceivers of the plurality as the patient moves throughout the medical facility.

9. The remote telemeter according to claim 8, wherein the transceiver circuit transmits like physiologic data to different transceivers during different timeslots to provide at least space and time diversity.

10. The remote telemeter according to claim 8, wherein the transceiver circuit transmits like physiologic data to different transceivers on different RF channels to provide at least frequency and spacial diversity.

11. The remote telemeter according to claim 8, wherein the transceiver circuit transmits a first packet to a first transceiver during a first timeslot on a first RF channel, and transmits a second packet to a second transceiver during a second timeslot on a second RF channel, wherein the first and second packets contain like physiologic data.

23

12. The remote telemeter according to claim 8, wherein the transceiver circuit transmits timeslot request messages to selected transceivers.

13. The remote telemeter according to claim 8, wherein the transceiver circuit receives signals transmitted by the transceivers, and the processor monitors signal strengths of said signals to select transceivers to use.

14. The remote telemeter according to claim 8, wherein the transceiver circuit receives patient location signals from the transceivers.

15. The remote telemeter according to claim 8, wherein the transceiver circuit and the processor are contained within a housing which is adapted to be attached to an ambulatory patient.

16. A method of transferring physiologic data from a wireless telemeter device which attaches to a patient to a monitoring station of a patient monitoring system, the method comprising:

receiving real-time physiologic data from at least one sensor which attaches to the patient;

transmitting the physiologic data from the telemeter device to a first transceiver of the patient monitoring system during a first timeslot on a first wireless channel; and

transmitting the physiologic data to a second transceiver of the patient monitoring system during a second timeslot on a second wireless channel, wherein the first and second transceivers are positioned remotely from one another to provide overlapping coverage zones.

17. The method according to claim 16, further comprising monitoring signals transmitted by a plurality of transceivers of the patient monitoring system, and using the signals to select subsets of transceivers with which to establish wireless communications links.

18. The method according to claim 16, further comprising receiving a timeslot assignment message from a transceiver of the patient monitoring system.

24

19. A method of transferring physiologic data from a wireless telemeter device which attaches to a patient, to a monitoring station of a patient monitoring system, the method comprising:

receiving real-time physiologic data from at least one sensor which attaches to the patient;

receiving signals on a plurality of wireless channels from a plurality of transceivers that are coupled to the monitoring station, the transceivers spatially distributed throughout the medical facility to provide overlapping zones of data reception coverage;

using said signals to periodically assess wireless link conditions provided by individual transceivers of the plurality; and

based on assessments of the wireless link conditions, selecting at least one transceiver to use to transfer the real-time physiologic data to the monitoring station;

wherein the method accommodates patient mobility by selecting a transceiver that corresponds to the patient's current location.

20. The method according to claim 19, further comprising transmitting like real-time physiologic data to multiple different transceivers on multiple different channels to provide at least space and frequency diversity.

21. The method according to claim 20, wherein transmitting like real-time physiologic data to multiple different transceivers comprises transmitting a unit of physiologic data to a first transceiver during a first timeslot and transmitting the unit of physiologic data to a second transceiver during a second timeslot, to thereby additionally provide spacial diversity.

22. The remote telemeter according to claim 1, wherein each of the plurality of wireless channels is a frequency division multiplexed channel.

* * * * *



US006154461A

United States Patent [19][11] **Patent Number:** 6,154,461

Sturniolo et al.

[45] **Date of Patent:** Nov. 28, 2000[54] **SEAMLESS ROAMING AMONG MULTIPLE NETWORKS**[75] **Inventors:** Emil A. Sturniolo, Medina; Frank D. Ciotti, Jr., Solon; Krishna P. Seshadri, Medina, all of Ohio[73] **Assignee:** Telxon Corporation, The Woodlands, Tex.[21] **Appl. No.:** 08/856,122[22] **Filed:** May 14, 1997[51] **Int. Cl.⁷** H04L 12/46[52] **U.S. Cl.** 370/401; 370/338[58] **Field of Search** 370/310, 328, 370/338, 401, 402, 403, 404, 349, 473; 455/432, 550, 556, 561, 575[56] **References Cited****U.S. PATENT DOCUMENTS**

4,837,800	6/1989	Freeburg et al. .	
4,912,756	3/1990	Hop .	
5,426,637	6/1995	Derby et al.	370/401
5,446,736	8/1995	Gleeson et al.	370/473
5,490,139	2/1996	Baker et al.	370/401
5,566,225	10/1996	Haas .	
5,572,528	11/1996	Shuen 370/402	
5,574,774	11/1996	Ahlberg et al. .	
5,664,007	9/1997	Samadi et al. .	
5,752,185	5/1998	Ahuja .	
5,796,727	8/1998	Harrison et al. .	
5,848,064	12/1998	Cowan 370/401	
5,856,974	1/1999	Gervais et al. 370/401	
5,889,816	3/1999	Agrawal et al. .	


OTHER PUBLICATIONS


Mobile Network Computing Protocol (MNCP), D. Piscitello, L. Phifer, Y. Wang, R. Hoyey, Aug. 28, 1997.

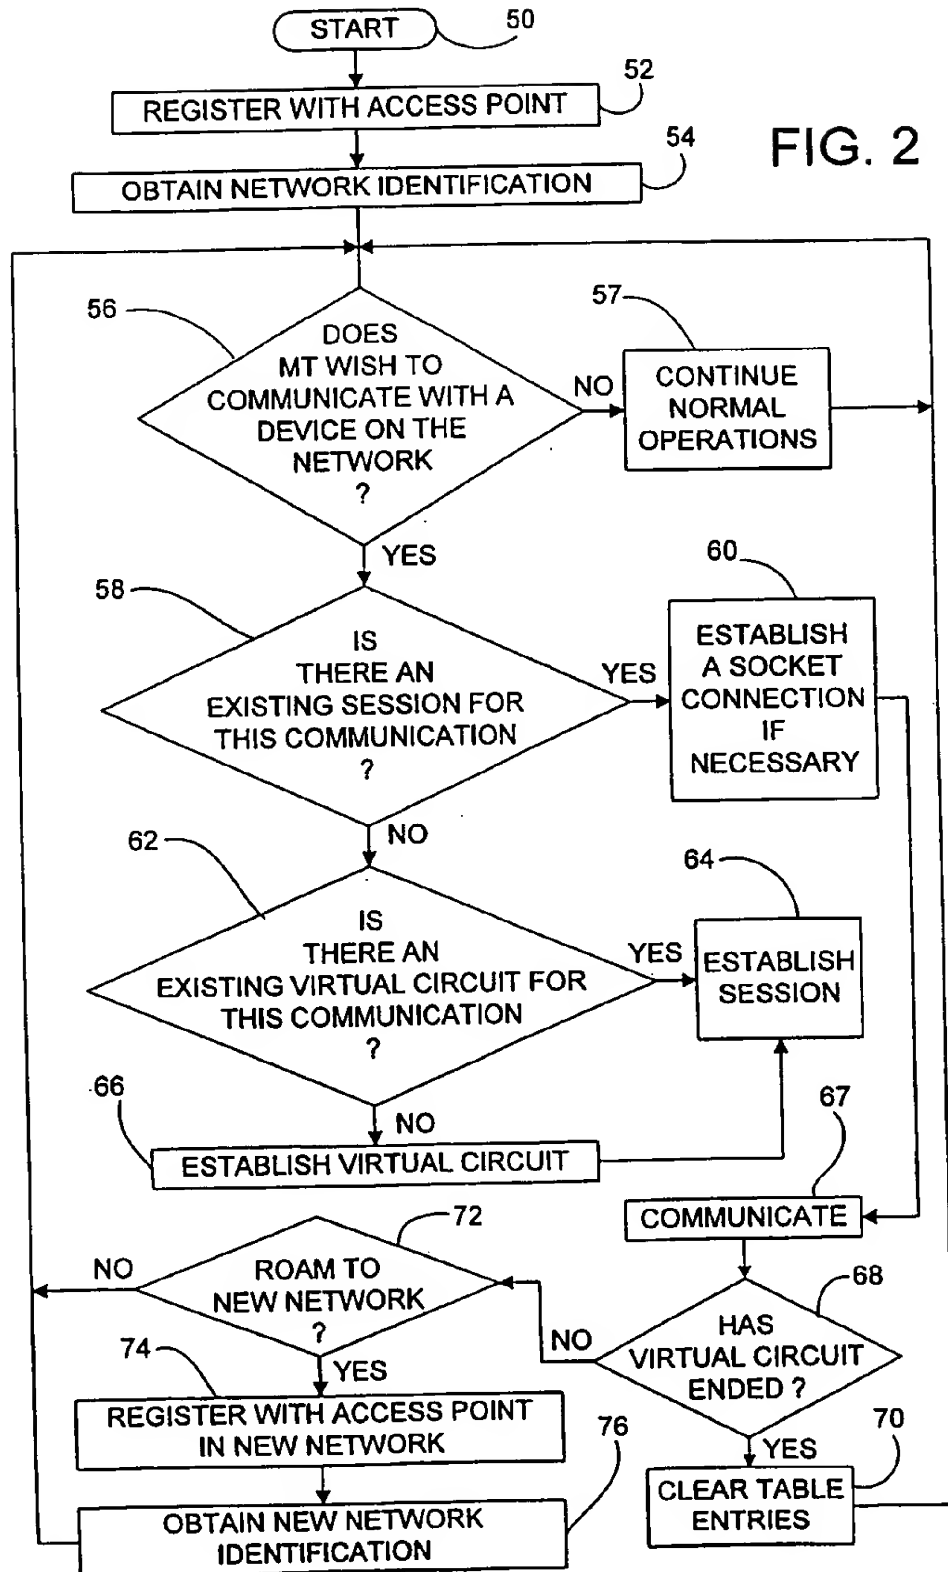
Primary Examiner—Chi H. Pham
Assistant Examiner—Frank Duong
Attorney, Agent, or Firm—Renner, Otto, Boisselle & Sklar, LLP





[57] **ABSTRACT**






The present invention relates to wireless communication systems involving multiple local area networks. A communication system according to the present invention includes a plurality of local area networks (LANs). Each of the LANs includes: a network backbone; and at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with a device on the network backbone via the at least one access point. When the mobile terminal is registered to at least one access point in one of the plurality of LANs the mobile terminal is assigned a first network address, and when the mobile terminal is registered to at least one access point in another of the plurality of LANs the mobile terminal is assigned a second network address in place of the first network address, the second network address being different from the first network address. The mobile communication system also includes a system backbone interconnecting the plurality of LANs for permitting communications between the plurality of LANs. Furthermore, the system includes a gateway controller, operatively coupled to one of the plurality of LANs, for serving as an intermediary for communications between the mobile terminal and a device on one of the system backbones in order that in the event the mobile terminal is assigned a different network address by virtue of registering with an access point in another of the LANs, the device is able to maintain communications with the mobile terminal without requiring knowledge of a change in the network address of the mobile terminal.


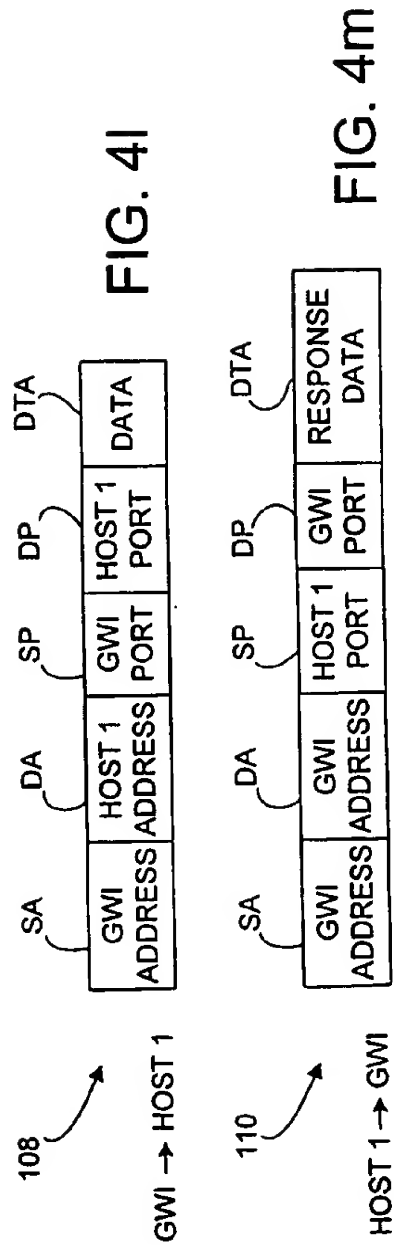
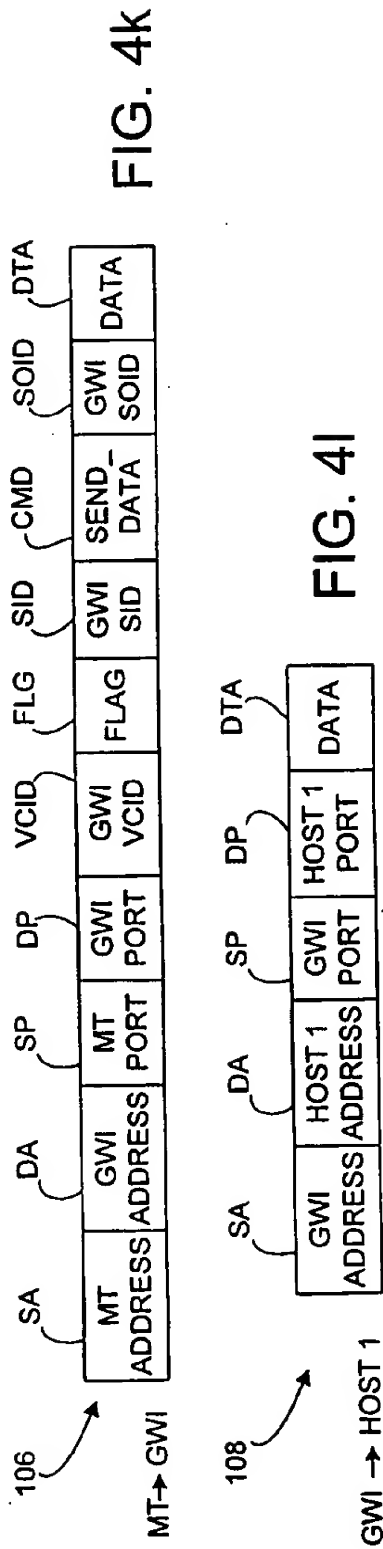


22 Claims, 9 Drawing Sheets















VIRTUAL CIRCUIT TABLE IN MT

		MT SESSION 1	GW SESSION 1	MT SOCKET 1	GW SOCKET 1
				MT SOCKET 2	GW SOCKET 2
		MT SESSION 2	GW SESSION 2	•	•
		• • •	• • •	•	•
		MT SESSION L	GW SESSION L	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•

FIG. 5a

VIRTUAL CIRCUIT TABLE IN GW

GW VC ₁	MT VC ₁	GW SESSION 1	MT SESSION 1	GW SOCKET 1	MT SOCKET 1
		GW SESSION 2	MT SESSION 2	GW SOCKET 2	MT SOCKET 2
		• • •	• • •	• • •	• • •
		GW SESSION L	MT SESSION L	•	•
GW VC ₂	MT VC ₂	•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
• • • •	• • • •	•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•
GW VC _z	MT VC _z	•	•	•	•
		•	•	•	•
		•	•	•	•
		•	•	•	•

FIG. 5b




FIG. 6




FIG. 7

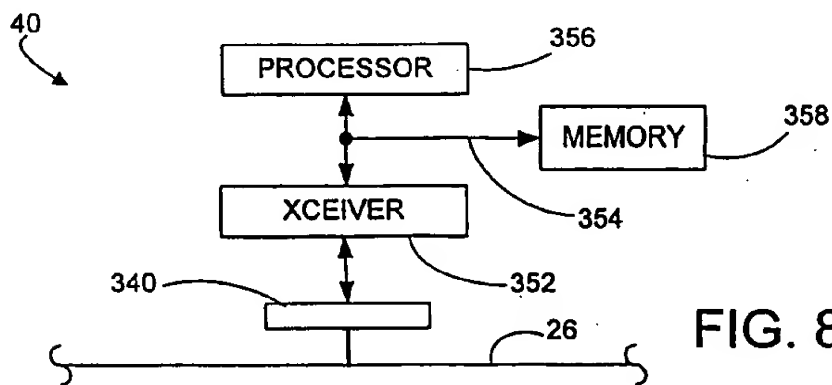


FIG. 8

SEAMLESS ROAMING AMONG MULTIPLE NETWORKS

TECHNICAL FIELD

The present invention relates generally to wireless mobile communication systems, and more particularly to mobile communication systems involving multiple local area networks (LANs).

BACKGROUND OF THE INVENTION

In recent years, the use of wireless mobile communication systems has become increasingly popular. For example, wireless mobile terminals now serve to help automate and expedite processes in retail, manufacturing, warehousing and other industries. In a retail environment, wireless mobile terminals may take the form of a wireless bar code reading device for use in tracking inventory and checking prices. In the warehousing industry, the same mobile terminals may be used to keep accurate accounting of incoming and outgoing shipments. In health care, transportation and other industries, the mobile terminals may take the form of wireless pen based computers to aid with on-site document control procedures, etc. In order to provide for real time communication, the mobile terminals often include a radio which allows them to communicate with a host computer connected to a LAN, for example.

LANs typically allow for connecting of devices operating in a building or specified site. Devices physically connected to the LAN may include desk top computers, printers and host computers. If the LAN also supports wireless mobile terminals such as those mentioned above, the LAN will also have connected thereto one or more access points (sometimes referred to as base stations). Each access point is coupled to the LAN and includes at least one radio through which wireless communication with the mobile terminals can occur.

Each access point can communicate with mobile terminals operating within the cell coverage area of the access point. The cell coverage area is the area in which the access point can reliably communicate with a mobile terminal. Once the mobile terminal roams outside of the cell coverage area of the access point, the mobile terminal can no longer communicate with the LAN through that particular access point. In order to provide cell coverage throughout an entire building or site, a LAN typically includes multiple access points strategically located throughout the building or site. Thus, the combined cell coverage of the access points is sufficient to cover the entire building or site. Mobile terminals may then roam from one area to another within the LAN.

As a mobile terminal roams throughout a LAN, it is important that an end to end session established between the mobile terminal and a device coupled to the backbone be maintained as the mobile terminals move from one cell to another cell. Known techniques for providing such seamless roaming from one access point to another access point within a given LAN are described in U.S. Pat. No. 5,276,680, for example.

However, recent trends have shown a desirability for mobile terminals to be able to roam not only within a given LAN, but also among different LANs and/or wide area networks (WANs). Although there are known techniques for allowing mobile terminals to roam seamlessly from one access point to another within a given LAN, this does not include the ability for mobile terminals to roam seamlessly from LAN to LAN, or LAN to WAN, for example. Accord-

ing to current technology, when a mobile terminal wishes to roam from one LAN to another the mobile terminal typically must disassociate itself from one LAN and reassociate itself with another LAN. This break in communication makes it difficult for information originating from the former LAN to be forwarded to the mobile terminal in the new LAN. Such difficulty adds overhead such as time and complexity associated with establishing a new connection.

Further, regardless of whether a mobile terminal roams to different LANs or WANs, information may also be lost if a connection (also referred to as a session) between the mobile terminal and some device on the LAN terminates for whatever reason. Such termination of a session may occur, for example, due to the mobile terminal not responding to communication from the device on the LAN due to the mobile terminal having roamed out of coverage area or because the mobile terminal is in a power savings mode.

In view of the aforementioned shortcomings with conventional techniques, it will be appreciated that there is a strong need in the art for a wireless mobile communication system which provides for seamless roaming among different networks. In particular, there is a strong need in the art for a system which provides for seamless roaming between different LANs each forming part of a WAN.

SUMMARY OF THE INVENTION

The communication system of the present invention introduces a gateway controller (hereinafter referred to simply as a "gateway") connected to at least one network such as a LAN or WAN. Each gateway functions as an intermediary for communications between mobile terminals registered to an access point within a network or otherwise coupled to the network and one or more other devices. By serving as an intermediary, the actual network addresses of the mobile terminals become transparent to the devices with which the mobile terminals are communicating. As a result, even if a mobile terminal roams from one LAN to another LAN and receives a new network address, communications between the mobile terminal and the other devices are not interrupted so as to provide for seamless roaming.

A primary difficulty in providing seamless roaming between different LANs and/or WANs had been the ability of the different networks to properly identify a mobile terminal as the same device. When the mobile terminal associates itself with a different network, a server on the network typically supplies the mobile terminal with a new network address (ID) from its list of available IDs. Since devices in the network with which the mobile terminal was previously communicating with ordinarily would not know the new network address of the mobile terminal, seamless roaming was difficult. With the gateway of the present invention, however, it is not necessary for the devices to know the new network address of the mobile terminal. The gateway routes the communications between the mobile terminals and the respective devices such that there is no need to first inform the devices of the new network address. Thus, by use of the gateway, seamless mobile terminal roaming is achieved.

According to one particular aspect of the invention, a communication system is provided, including: a plurality of local area networks (LANs), each of the LANs including: a network backbone; and at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with a device on the network backbone via the at least one access point; wherein when the

mobile terminal is registered to at least one access point in one of the plurality of LANs the mobile terminal is assigned a first network address, and when the mobile terminal is registered to at least one access point in another of the plurality of LANs the mobile terminal is assigned a second network address in place of the first network address, the second network address being different from the first network address; a system backbone interconnecting the plurality of LANs for permitting communications between the plurality of LANs; and a gateway controller, operatively coupled to one of the plurality of LANs, for serving as an intermediary for communications between the mobile terminal and a device on one of the system backbones in order that in the event the mobile terminal is assigned a different network address by virtue of registering with an access point in another of the LANs, the device is able to maintain communications with the mobile terminal without requiring knowledge of a change in the network address of the mobile terminal.

In accordance with another aspect of the invention, a gateway controller for use in a mobile communication system including a plurality of local area networks (LANs), each of the LANs including a network backbone; and at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with a device on the network backbone via the at least one access point is provided, the gateway controller including: means for operatively coupling the gateway controller to the network backbone of a selected one of the plurality of LANs; a gateway address circuit for communicating with a mobile terminal registered with the access point of the selected LAN via the network backbone for establishing indicia distinguishing the mobile terminal from among other mobile terminals; a communication circuit for receiving communications from the registered mobile terminal intended for the device and forwarding the received communications to the device together with at least part of the indicia as provided by the communication circuit, and for receiving communications, including the at least part of the indicia, from the device intended for the registered mobile terminal and forwarding the received communications to the registered mobile terminal based on the at least part of the indicia.

According to yet another aspect of the invention, a network communication system is provided, including: a plurality of local area networks (LANs), each of the LANs including: a network backbone; and at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with the network backbone via the at least one access point; a system backbone interconnecting the plurality of LANs for permitting communication between the plurality of LANs; and a gateway controller, operatively coupled to one of the plurality of LANs and able to communicate with devices on other of the plurality of LANs via the system backbone, the gateway controller serving as an intermediary for communications between the mobile terminal and a device on a first of the plurality of LANs; the gateway controller assigning a gateway address to the mobile terminal such that once the mobile terminal starts a session with the device coupled to the first of the plurality of network backbones, the mobile terminal may continuously maintain the session through the gateway even if the mobile terminal registers with an access point coupled to another of the plurality of network backbones.

To the accomplishment of the foregoing and related ends, the invention, then, comprises the features hereinafter fully

described and particularly pointed out in the claims. The following description and the annexed drawings set forth in detail certain illustrative embodiments of the invention. These embodiments are indicative, however, of but a few of the various ways in which the principles of the invention may be employed. Other objects, advantages and novel features of the invention will become apparent from the following detailed description of the invention when considered in conjunction with the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a wireless mobile communication system in accordance with the present invention;

FIG. 2 is a flowchart describing, in relevant part, the general operation of the mobile communication system in accordance with the present invention;

FIG. 3 is a schematic diagram showing a general format of an information packet in accordance with the present invention;

FIGS. 4a-4j are schematic illustrations representing an exchange of packets for establishing a virtual circuit, session and socket connection between a mobile terminal and another device on the network in accordance with the present invention;

FIGS. 4k-4o are schematic illustrations representing an exchange of data between a mobile terminal and another device before and after the mobile terminal moves from one LAN to another LAN;

FIG. 5a represents a virtual circuit table maintained in memory by each mobile terminal for keeping track of the various virtual circuits which have been established through a gateway with different devices in accordance with the present invention;

FIG. 5b represents a corresponding virtual circuit table maintained in memory by each gateway for keeping track of the various virtual circuits which have been established through the gateway between various mobile terminals and devices in accordance with the present invention;

FIG. 6 is a block diagram of a mobile terminal in accordance with the present invention;

FIG. 7 is a block diagram of an access point in accordance with the present invention; and

FIG. 8 is a block diagram of a gateway in accordance with the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention will now be described with reference to the drawings wherein like reference numerals are used to refer to like elements throughout.

As is mentioned above, the present invention relates to wireless communication systems which include mobile terminals that can roam from cell to cell and from LAN to LAN. Such mobile terminals can be data terminals, telephones, pagers, etc. In the exemplary embodiment described hereinafter, each mobile device is a mobile data terminal (hereinafter "mobile terminal") used to communicate data such as inventory or the like within a communications system such as a cellular communication system. However, it is recognized that the invention contemplates other types of mobile terminals and is not intended to be limited to systems utilizing mobile data terminals. Other types of mobile terminals may be referred to in the industry as a mobile end system, mobile node, or mobile client, for example.

Referring now to FIG. 1, a communication system 20 is shown in accordance with the exemplary embodiment of the present invention. The communication system network 20 includes a plurality of LANs (e.g., LAN1-LAN3) each coupled together via a network backbone 26. Each LAN1-LAN3 itself forms a communication network. The LANs are interconnected according to generally known network principles by way of a system backbone 24, and specifically in the present embodiment by a WAN system backbone 24. It shall be appreciated, however, that the system backbone 24 need not be wireless in nature but rather hardwired such as those achieved by connecting to an intranet or internet, for example, which could also serve as the system backbone 24.

Each of the LANs (LAN1-LAN3) have generally the same configuration, hence only LAN1 will be described in detail. However, it will be appreciated that there may be variations in the respective LANs without departing from the scope of the invention. Referring initially to LAN1, the local area network comprises its own network backbone 26. The network backbone 26 may be a hardwired data communication path made of twisted pair cable, shielded coaxial cable or fiber optic cable, for example, or may be wireless in nature. Connected to the network backbone 26 are several access points 28, only one of which is shown (namely, access point AP1) for sake of illustration. Each access point 28 serves as a point through which wireless communications may occur with the network backbone 26. Additionally, in order to expand the effective communication range of the access points 28, one or more wireless access points (not shown) also may be included in LAN1.

Each access point 28 is capable of wirelessly communicating with other devices within its cell coverage via an antenna 32. As is known, depending on the type of antenna 32 selected and the output power of the respective access point 28 the cell coverage of the access point 28 may take any of several different forms and sizes. For example, FIG. 1 depicts the access point 28 as utilizing an omni-directional antenna 32 wherein a generally spherical cell area of coverage is obtained. However, a directed yagi-type antenna or other form of antenna could also be used as will be readily appreciated.

The LAN1 also includes one or more mobile terminals 36. For sake of example, only one mobile terminal 36 is shown although it will be appreciated that each LAN is likely to have several mobile terminals 36 associated therewith. Each mobile terminal 36 communicates with devices on the network backbone 26 of the LAN in which it is registered or, as described below, is capable of communicating with devices on the network backbone 26 of other LANs within the WAN. Within a given LAN (e.g., LAN1), the mobile terminal 36 may roam from one cell to another as covered by different access points 28. While roaming within a given LAN, the mobile terminal 36 is configured to associate itself with a new access point 28 in each new cell area according to conventional techniques. The mobile terminal 36 communicates with the network backbone via wireless communications through the access point 28.

According to the present invention, however, mobile terminals 36 also can seamlessly roam from one network to another network without a need to terminate and reestablish an end to end session between the mobile terminal 36 and a device coupled to one of the networks. For example, FIG. 1 illustrates in phantom the manner in which a mobile terminal 36 roams from LAN1 to LAN2. Specifically, the mobile terminal 36 originally is registered to an access point 28 (AP1) in LAN1. The mobile terminal 36 originally has a

network identification or address which has been assigned to it by virtue of being registered within LAN1. However, when the mobile terminal 36 moves outside of the cell coverage of access point AP1 and into the cell coverage of an access point 28 (AP2) included in LAN2, the mobile terminal 36 newly registers with the access point AP2. As a result, the mobile terminal 36 receives a new network identification or address by virtue of becoming registered within LAN2.

Since the devices which the mobile terminal 36 had been communicating with in LAN1 would not know the new address assigned to the mobile terminal 36 upon moving to LAN2, the present invention provides a means for compensating for such lack of knowledge. Specifically, in the present embodiment each LAN also includes a gateway 40 which serves as an intermediary for communications between the mobile terminal 36 and other devices within the system 20. However, it will be appreciated that only one gateway 40 need be connected to one of the network backbones 26 to carry out the present invention. Mobile terminals registered with access points 28 on a network backbone 26 other than the network backbone 26 the gateway 40 is connected to would communicate with the gateway 40 via the system backbone 24 discussed above. As is described more fully below, for each mobile terminal 36, a virtual circuit is established between itself and the gateway 40. Although the network address of the mobile terminal 36 may change as a result of the mobile terminal 36 roaming from one LAN to another LAN, the relevant parameters of the corresponding virtual circuit remain the same. Hence, communications between the mobile terminal 36 and a given device are properly routed notwithstanding the change in the network address of the mobile terminal 36. In fact, in the preferred embodiment the devices with which the mobile terminal 36 is communicating remain unaware that the mobile terminal 36 has received a new network address. The manner in which such seamless roaming is carried out is described in more detail below in relation to FIGS. 2-5.

Also connected to the network backbone 26 of each LAN is a host computer 42 and a domain name server 44. The host computer 42 performs conventional host functions within the respective LAN. In addition, the host computer 42 serves as an interface connection to the system backbone 24 in a conventional manner. Of course, the gateway 40 or any other LAN device could alternatively serve as an interface connection to the system backbone 24. The domain name server 44 in each LAN performs the conventional function of providing a name to network address mapping for devices within each LAN. Furthermore, each LAN may include one or more other devices connected to the network backbone 26. Although not shown, such devices may include work terminals, printers, cash registers, etc.

Turning now to FIG. 2, the basic operating protocol for a mobile terminal 36 roaming between LANs is shown in accordance with the invention. For sake of example, it is assumed that the mobile terminal 36 initially powers up and registers within the cell area of an access point 28 (AP1) belonging to LAN1. Thereafter, the mobile terminal 36 will move to LAN2 as represented in FIG. 1 and register with access point AP2 in LAN2. It will be appreciated, however, that the same principles are applied when roaming between any two LANs.

Referring initially to step 50, the mobile terminal 36 is powered up and/or reset within the cell coverage of access point AP1. Next, in step 52 the mobile terminal 36 registers with the access point AP1 using any of several known conventional techniques, for example. By registering, the

access point AP1 assumes responsibility for receiving wireless communications from the mobile terminal 36 and forwarding the communications onto the network backbone 26. Similarly, the access point AP1 assumes responsibility for receiving communications on the network backbone 26 which are destined for the mobile terminal 36. The access point AP1 then forwards such communications wirelessly to the mobile terminal 36.

Next, in step 54 the mobile terminal 36 obtains a network identification (ID)/address. Such network ID may be obtained via any of several known conventional techniques in which a unique network ID is assigned to each particular mobile terminal 36 within the LAN1. In some instances, the network ID and/or port ID may be statically preconfigured within the mobile terminal 36. In the exemplary embodiment, the mobile terminal 36 in step 54 obtains a network ID which includes a network address and port in accordance with conventional network protocol(s). The mobile terminal 36 also obtains the gateway 40 network ID and port ID in a similar fashion. Alternatively, the mobile terminal 36 may first obtain a link layer ID in accordance with the technique described in commonly assigned, co-pending U.S. patent application Ser. No. 08/778,405, filed Jan. 2, 1997 and entitled "Mobile Device ID Allocation System and Method" which in turn could be used to obtain a network ID. The disclosure of the '405 application is incorporated herein by reference.

Following step 54, the mobile terminal 36 in step 56 determines if any application program running on the mobile terminal 36 is currently attempting to communicate information to a device on the network backbone 26. If the mobile terminal 36 determines that no information currently needs to be transmitted, the mobile terminal goes to step 57 where it continues all of its other normal operations and returns again to step 56. If, however, the mobile terminal 36 does wish to communicate information to a device on the network, the mobile terminal 36 proceeds to step 58.

In step 58, the mobile terminal 36 determines if there is an existing session for the mobile terminal 36 to transmit the information it desires. As discussed in more detail below, the mobile terminal 36 may have already established a session with GATEWAY1 to communicate information to other devices throughout the communication system 20. If the mobile terminal 36 determines that a session is established for this information it desires to transmit then the mobile terminal 36 continues to step 60. If not, the mobile terminal 36 continues to step 62 where the mobile terminal determines if a virtual circuit has been established for transmitting the information.

Prior to being able to transmit the information to a device on the network, the mobile terminal 36 must, in step 60, determine if a socket end point is established (via GATEWAY1) to the device the mobile terminal 36 desires to communicate. Socket end points are frequently established and ended between a mobile terminal 36 and a device on the network and it may even be the case that more than one socket end point is established between a given mobile terminal 36 and a given device to accommodate the transfer of data between application programs running on each. Thus, in step 60, if a socket end point which is able to handle the transfer of current information from the mobile terminal 36 to the device is not established, the mobile terminal establishes such a socket end point. The socket end point is established based on an exchange of a series of packets between the mobile terminal 36 and GATEWAY1. This process is described below in connection with FIGS. 4g-4j. Once the socket is established, communications between the

mobile terminal 36 and the device occur in step 67. The communications occur via GATEWAY1 serving as an intermediary. An example of such communication is described below in association with FIGS. 4k-4n. It is noted that in this example, the mobile terminal 36 initiated a connection oriented socket end point with the device it desired to communicate, however, connectionless oriented socket end point associations could also be established using conventionally known protocols such as UDP.

If in step 58, the mobile terminal 36 determines that there is not an existing session to transfer the current information to the device the mobile terminal 36 will have to establish such a session. However, prior to establishing such a session, the mobile terminal 36 in step 62 determines if a virtual circuit exists between a mobile terminal 36 and one or more other devices within the communication system 20. As will be described below, each virtual circuit in the preferred embodiment can include up to a predefined number of sessions. Each session, in turn, consists of up to a predefined number of sockets. The predefined number of session and sockets available are each typically of a very large magnitude (i.e. 2^{32}) and therefore are often considered unlimited. If a virtual circuit already exists for transferring the current information as determined in step 62, the mobile terminal 36 goes to step 64 where it will establish a session within this virtual circuit prior to transferring the information to the device. If a virtual circuit does not exist, the mobile terminal 36 will have to first establish the virtual circuit before establishing the session as discussed below with respect to step 66.

Turning to step 64, a session is established in accordance with the present invention by exchanging a series of information packets between the mobile terminal 36 and GATEWAY1 as discussed below in relation to FIGS. 4c-4d and 5a-5b. If necessary address information for the particular device with which the mobile terminal wishes to communicate may be obtained using well known name resolution techniques. For example, the mobile terminal 36 may want to communicate with the host computer 42 (HOST1) in LAN1. As discussed below in connection with FIGS. 4e-4f, the mobile terminal 36 and GATEWAY1 exchange a series of packets which results in GATEWAY1 providing the mobile terminal 36 with the network address of the host computer HOST1. Once the session is established and the address of the particular device is determined, the mobile terminal continues to step 60 where a socket end is established.

If in step 62, the mobile terminal 36 determines that no virtual circuit exists for transferring the present information to the device, or if the mobile terminal 36 simply wishes to establish a new virtual circuit, the mobile terminal 36 goes to step 66. In step 66, a virtual circuit between the mobile terminal 36 and GATEWAY1 is established. As described below with respect to FIGS. 4a-45 and 5a-5b, the mobile terminal and GATEWAY1 exchange a series of information packets which establishes a virtual circuit entry in corresponding tables stored therein. The virtual circuit is used to identify a particular communication connection between the mobile terminal 36 and GATEWAY1 through which connections may be established with devices which the mobile terminal 36 wishes to communicate. Following the establishment of a virtual circuit in step 66, the mobile terminal continues to steps 64 and 60 where a session and socket are next established, respectively, within the virtual circuit.

Once a virtual circuit, session, and socket end point is established and the information is transmitted in step 67, the mobile terminal goes to step 68 to determine if any virtual circuits are to be terminated.

For example, the mobile terminal 36 may conclude that it no longer needs to communicate with the host computer HOST1 and wishes to terminate the connection. Alternatively, there may be a predefined time limit placed on the duration of each virtual circuit and thus all corresponding sessions. Accordingly, the GATEWAY1 may be responsible for serving as a timekeeper and determining when a virtual circuit is to terminate. The GATEWAY1 in such instance is responsible for informing the mobile terminal 36 that the virtual circuit is about to be terminated. If a virtual circuit is to end as determined in step 68, the system proceeds to step 70 in which the corresponding session and socket connection entries in the tables of the mobile terminal 36 and the GATEWAY1 are cleared as will be better understood in view of the discussion of FIGS. 5a-5b below.

If in step 68 the session has not ended, the system proceeds to step 72 in which the mobile terminal 36 determines it has moved to a new LAN (e.g., LAN2). The manner in which the mobile terminal 36 determines if it has moved to a new LAN can be based on known conventional techniques for determining entry to a new LAN. For example, such determination can be based on the known mobile internet protocol defined in Internet Engineering Task Force (IETF) Spec. R.F.C. 2002. Generally speaking, the mobile terminal 36 upon roaming from LAN1 to LAN2 will proceed beyond the cell coverage of AP1 in LAN1 and will not be able to register with any other access points 28 within LAN1. Hence, the mobile terminal 36 will conclude that it has roamed from the previous LAN (LAN1) to a new LAN (e.g., LAN2) with which it must newly register. It is noted that the steps of determining if a virtual circuit has ended (step 68) and the step of determining if the mobile terminal 36 roamed to a new LAN (step 72) may occur at any instant throughout the process described in FIG. 2 and is only shown at its current locations for discussion purposes only.

If the mobile terminal 36 concludes it has roamed to a new LAN, the mobile terminal 36 proceeds to step 74 in which it registers with an access point 28 (e.g., AP2) in the new LAN (e.g., LAN2) in the same manner as was done in step 52. Assuming the mobile terminal 36 has roamed within the cell coverage of the access point AP2, the mobile terminal 36 will thus be able to register with the access point AP2. Next, in step 76 the mobile terminal 36 obtains a new network ID pertaining to the new LAN2. Step 76 is similar to step 54 described above in that the mobile terminal 36 uses conventional techniques to obtain a unique network ID within its local network LAN2. There may or may not be communication between the different LANs 1-3 regarding the particular network IDs assigned to the mobile terminals 36 and the network IDs of the mobile terminals 36 are not fixed in the preferred embodiment. Thus, the new network ID of the mobile terminal 36 in step 76 may be different from the network ID assigned in step 54 in virtually every case.

Nevertheless, the virtual circuit, session, and socket connection information previously obtained via the GATEWAY1 is still valid despite the mobile terminal 36 receiving a new network address in step 76. Thus, packets delivered to the GATEWAY1 for routing to the mobile terminal 36 still may be routed to the mobile terminal 36 by the GATEWAY1 as discussed below in connection with FIG. 4a. Accordingly, the system returns to step 66 in which communications can be continued to be carried out seamlessly. If, on the other hand, the mobile terminal 36 does not roam to a new LAN as determined in step 72, the system returns directly to step 56. It is noted that under certain circumstances it may be the case that the mobile terminal 36 ends its session with a host computer or other device without the device knowing the

session has ended and therefore communication destined for the mobile terminal 36 in such cases would not be forwarded on the GATEWAY1.

Turning now to FIG. 3, a standard packet format for information communicated in the system 20 is shown. The packet format shown in FIG. 3 includes a number of different fields, some or all of which may be present in a given packet at a given time. The packet consists primarily of a header portion 80 and a payload portion 82. The header portion 80 is included in all communications in the respective LANs. The payload portion 82 ordinarily is reserved for data being transferred between two or more devices communicating in a given LAN. According to the present invention, however, the payload portion 82 also is used to communicate information regarding the particular virtual circuit, session, socket, etc. As will be appreciated, this allows the present invention to be carried out in a manner which is transparent to most devices on the network with the exception of the gateway 40 and mobile terminals 36. Thus, the present invention can be incorporated into existing systems with only minor modifications.

As shown in FIG. 3, each packet includes a source address (SA) field which identifies the network address of the device transmitting the packet. In addition, each packet includes a destination address (DA) field which identifies the network address of the device to which the packet is being transmitted. Next, the packet includes a source port (SP) field and destination port (DP) field which identify the ports of the devices transmitting and receiving the packet, respectively. Within the payload portion 80, the packet may include a globally unique identification (GUID) field used for uniquely identifying a GUID associated with the mobile terminal 36 transmitting the packet. A virtual circuit identification (VCID) field is used for identifying a particular virtual circuit associated with the packet transmission. A flag (FLG) field is used to indicate a change in the network identification of the mobile terminal 36 as will be described below. Similarly, a session identification (SID) field and socket identification (SOID) field are used for identifying the particular session and socket, respectively, associated with the packet transmission as will be explained below.

A command (CMD) field is included in the packet for identifying particular functions to be carried out as will be described by way of example below. A data (DTA) field is used for containing data which is to be transmitted in accordance with the invention.

According to conventional communication network routing protocol, information packets ultimately are routed to the destination address and destination port identified in the DA and DP fields, respectively. With regard to packets which are transmitted from the mobile terminal 36, each access point 28 is configured to recognize a packet which is being transmitted from a mobile terminal 36 which is registered thereto. The access point 28 in turn receives the packet via its antenna 32 and retransmits its contents onto the network backbone 26 according to known convention. As far as packets which are destined for the mobile terminal 36, the access point 28 receives such a packet off the network backbone 26 and retransmits the information to the mobile terminal 36 via its antenna 32 according to known convention.

With the exception of the mobile terminals 36, the locations of all fixed devices on the respective LANs 1-3 and WAN 22 are essentially known based on conventional network routing protocol. In other words, provided a current network identification is known for a device, a packet will

be delivered to the device via appropriate routers, etc. Thus, for example, the access point AP2 in LAN2 can transmit a packet to GATEWAY1 in LAN1 via the WAN backbone 24 according to conventional techniques. Hence, details regarding the communications between the fixed devices on the respective LANs will not be described in detail for sake of brevity.

Referring now to FIGS. 4a-4b, the manner in which a virtual circuit is established (FIG. 2, step 56) will be described. In the present example, the mobile terminal 36 initially registers to the access point AP1 in LAN1 as shown in FIG. 1. Next, as shown in FIG. 4a, the newly registered mobile terminal 36 generates and wirelessly transmits a packet 86 to the local gateway 40 (GATEWAY1). Using conventional techniques, the mobile terminal 36 retrieves its network ID and that of the gateway's 40 if such IDs are not already statically programmed into the mobile terminal 36. The SA field and SP field in packet 86 include the network address and port of the mobile terminal 36, respectively. The DA field and DP field include the network address and port of the GATEWAY1, respectively. The network address and port of the GATEWAY1 may be previously obtained by the mobile terminal 36 in a number of different manners. For example, each access point 28 may be programmed to provide the address and port for the default gateway 40 whenever a mobile terminal 36 newly registers. Alternatively, the network address and port for a gateway 40 may be statically configured within the mobile terminal 36.

The GUID field in the packet 86 is set to all zeros or some other predefined value to indicate that a GUID has not yet been assigned to the mobile terminal 36. Of course, in an alternative embodiment, the GUID may have been previously retrieved or preassigned by the mobile terminal 36. As is described below in relation to FIG. 4b, this prompts the GATEWAY1 to perform a function call to the LAN1 operating system requesting a GUID for the mobile terminal 36.

Referring briefly to FIG. 5a, shown is a virtual circuit table which is maintained in memory by the mobile terminal 36 for purposes of keeping track of the various virtual circuit connections. In the exemplary embodiment, the table consists of N rows with each row representing a possible virtual circuit (i.e., VC1 through VCN). N can be any preselected integer which represents the number of virtual circuits the mobile terminal 36 may have established at any given time. For each virtual circuit there is also a corresponding gateway 40 identification associated therewith (not shown) given that a mobile terminal could begin communicating via two or more gateways at any given time.

As shown in FIG. 5a, each virtual circuit entry (VC1 through VCN) has associated therewith L possible sessions (i.e., Session 1 through SessionL). L can be any preselected integer which represents the number of sessions in which the mobile terminal 36 can participate at any given time. Each session entry in the table has associated therewith a preselected number of possible sockets for representing socket end points in a given session. In the exemplary embodiment, each session includes two sockets (Socket1, Socket2), although a different number is possible. For each socket end point there is associated a device (i.e., HOST1) with which the mobile terminal 36 may communicate. As there may be many application programs running on the mobile terminal, there may also be several socket end points for the same device.

The mobile terminal virtual circuit table is empty initially and void of any valid columns or rows. As the mobile terminal 36 establishes new virtual circuits, sessions and

sockets to establish communication links, entries to the table are inserted with corresponding information from the gateway 40 to indicate which virtual circuits, sessions and sockets are utilized. Thus, the virtual circuit table can dynamically grow in size.

When a particular virtual circuit, session and/or socket is not in use and/or is terminated, the mobile terminal 36 invalidates or deletes the corresponding entries in the virtual circuit table thereby possibly reducing it in size. Thus, the mobile terminal 36 may always look to the virtual circuit table to see which circuits, sessions and sockets are available and which are in use. It is noted that virtual circuit table could be configured to accept only a limited number of entries or that the virtual circuit table could be static in nature in alternative embodiments.

Turning briefly to FIG. 5b, a corresponding virtual circuit table which is stored in memory in each gateway 40 is shown. The table is structured in the same manner as the virtual circuit table described in FIG. 5a and may be static or dynamic in nature. The number of virtual circuits will typically be larger than those found in the mobile terminal 36. This is because each gateway 40 will likely be handling traffic for multiple mobile terminals 36. The number of sessions per virtual circuit and sockets per session preferably is the same as that for the mobile terminal 36. Hence, a given virtual circuit connection between the gateway 40 and the mobile terminal can provide for L sessions with two or more sockets per session, for example. Also, for each socket the table includes information pertaining to the particular mobile terminal as is discussed below.

Like the mobile terminal 36, the gateway 40 clears any entries which are not in use. Thus, the gateway 40 can look to its virtual circuit table at any time and determine which particular connections are in use and which are available. Although the present embodiment shows all entries related to the virtual circuit, sessions and sockets to be included in a respective single table such as those shown in FIG. 5a and 5b, it will be appreciated that separate tables could be maintained for each and indexed appropriately to provide association.

Referring then back to FIG. 4a, the VCID field in the packet 86 includes the identification of a mobile terminal virtual circuit which the mobile terminal 36 preselects for purposes of establishing a virtual circuit with the GATEWAY1. Specifically, the mobile terminal 36 is configured to look to its virtual circuit table (FIG. 5a) and select a particular virtual circuit (e.g., VC1-VCN) which is presently available. The mobile terminal 36 selects one of the virtual circuits (nominally labeled MT VCID) and includes it in the VCID field. This serves to inform the GATEWAY1 of the particular virtual circuit label the mobile terminal 36 is using to represent the connection which is to be established. Finally, the CMD field includes a predefined function call "VCKT_START" intended to inform the GATEWAY1 that the mobile terminal 36 wishes to establish a virtual circuit.

The mobile terminal 36 then transmits the packet 86 wirelessly to the GATEWAY1 via the access point AP1. The GATEWAY1 receives the packet 86 and processes the packet as follows. The GATEWAY1 recognizes the VCKT_START command and looks to its virtual circuit table (FIG. 5b) in order to select a particular virtual circuit (e.g., VC1-VCZ) which is presently available. The GATEWAY1 selects a particular virtual circuit (e.g., VC1) based on what is available (i.e., not in use). The GATEWAY1 then stores in its table in association with the selected virtual circuit entry the MT VCID obtained from the VCID field of the packet

86. Thus, the selected virtual circuit thereby becomes associated with the virtual circuit identified by MT VCID.

Moreover, the GATEWAY1 is configured to detect the presence of all zeros in the GUID field. In response, the GATEWAY1 performs a function call to the operating system of the LAN1 to obtain a GUID for the mobile terminal 36. The manner in which the operating system can provide such GUID is known in the art.

In response to the packet 86, the GATEWAY1 generates a packet 88 as shown in FIG. 4b. The purpose of the packet 88 is to inform the mobile terminal 36 that a virtual circuit has been established and to notify the mobile terminal 36 of the particular virtual circuit label being used to define the connection by the GATEWAY1. In addition, the packet 88 is used to inform the mobile terminal 36 of the particular GUID which is being assigned thereto. Thus, the SA and SP fields of the packet 88 correspond to the address and port of the GATEWAY1, respectively. The DA and DP fields of the packet 88 correspond to the address and port of the mobile terminal 36, respectively, as obtained from the SA and SP fields of packet 86. The GUID field contains the GUID for the mobile terminal 36 (MT GUID) as obtained from the operating system. The VCID field contains the virtual circuit selected by the GATEWAY1 (nominally labeled GW1 VCID). The CMD field contains the predefined command "VCKT_STARTED" to indicate to the mobile terminal 36 that the virtual circuit has been started.

The GATEWAY1 then proceeds to transmit the packet 88 to the mobile terminal 36 via the access point AP1. The mobile terminal 36 receives the packet 88 and processes the packet 88 by storing in memory the MT GUID. In addition, the mobile terminal 36 stores in its virtual table entry for MT VCID the corresponding GW1 VCID as obtained from the VCID field of the packet 88.

Next, when the mobile terminal 36 is prepared to start a session (FIG. 2, step 64) the mobile terminal 36 generates a packet 90 to be sent to the GATEWAY1 as represented in FIG. 4c. The header portion of the packet 90 is identical to that of packet 86 (FIG. 4a). The VCID field includes the GW1 VCID assigned to the virtual circuit by the GATEWAY1 (as obtained from the virtual circuit table of the mobile terminal 36). In the SID field the mobile terminal 36 includes the identity of a session which the mobile terminal 36 selects for purposes of establishing the session. Specifically, the mobile terminal 36 again refers to its virtual circuit table (FIG. 5a) and selects a particular session (e.g., Session1) which is available within the previously selected virtual circuit (e.g., VCI). The mobile terminal 36 includes the session (nominally labeled MT SID) in the SID field of the packet 90. Finally, the mobile terminal 36 includes the predefined command "START_SESSION" in the CMD field to notify the GATEWAY1 that the mobile terminal 36 wishes to start a session. The mobile terminal 36 then proceeds to transmit the packet 90 to the GATEWAY1.

The GATEWAY1 receives the packet 90. In response to the command START_SESSION, the GATEWAY1 is configured to again look to its virtual circuit table (FIG. 5b) for the virtual circuit corresponding to the GW1 VCID identified in the VCID field. Specifically, the GATEWAY1 selects a particular session which is available in conjunction with the GW1 VCID previously selected for purposes of generating the packet 88. The GATEWAY1 then proceeds to store in the table entry corresponding to the selected session the session identification (MT SID) selected by the mobile terminal 36 as provided in the SID field of the packet 90.

The GATEWAY1 then proceeds to generate a response packet 92 as shown in FIG. 4d. The purpose of the response

packet 92 is to inform the mobile terminal 36 of the particular session label which the GATEWAY1 is assigning to the virtual circuit connection. Hence, the header portion is again standard and is identical to the header portion included in the packet 88. The VCID field includes the MT VCID for the virtual circuit as obtained from the virtual circuit table of the GATEWAY1 corresponding to the entry for GW1 VCID. The SID field includes the particular session (nominally labeled GW1 SID) selected by the GATEWAY1 in response to the packet 90. The CMD field includes a predefined command "SESSION_STARTED" indicating to the mobile terminal 36 that the gateway has selected a corresponding session. The GATEWAY1 then proceeds to transmit the packet 92 to the mobile terminal 36.

The mobile terminal 36 receives the packet 92 and looks to the entry in its virtual circuit table corresponding to the MT VCID identified in the VCID field of the packet 92. The mobile terminal 36 then proceeds to store in its virtual circuit table the GW1 SID obtained from the SID field of the packet 92 in association with the particular session identified in the SID field of the packet 90. Thus, a virtual circuit and session are established between the mobile terminal 36 and the GATEWAY1 based on an exchange of the packets represented in FIGS. 4a-4d. The mobile terminal 36 and the GATEWAY1 each know the corresponding VCID and SID for the other device.

As referred to in step 64 of FIG. 2, the mobile terminal 36 will need to obtain network addressing information for the device on the network with which it desires to communicate in order to establish a socket connection with a particular device on the network. In the present example, it is assumed that the mobile terminal 36 wishes to communicate with the host computer HOST1 in the LAN1. The mobile terminal 36 is configured to first generate a packet 94 as shown in FIG. 4e for the purpose of requesting address information. The packet 94 is to be transmitted to the GATEWAY1 and hence has a header portion similar to that described above in connection with the packet 86. The VCID and SID fields include the GW1 VCID and GW1 SID information, respectively, corresponding to the particular virtual circuit and session the requested address information will relate to. The CMD field includes a predefined command "GET_HOST_BY_NAME" or some other command identifying the particular device with which the mobile terminal 36 wishes to communicate. The DTA field of the packet 94 includes the particular name of the device (e.g., HOST1 NAME).

The packet 94 is transmitted to the GATEWAY1 which is configured to retrieve the device name from the DTA field and query the local name resolver such as DNS 44 for the corresponding network address. Next, the GATEWAY1 generates a response packet 96 as shown in FIG. 4f which is to be transmitted to the mobile terminal 36 with the appropriate header portion. The packet 96 includes the MT VCID and MT SID in the VCID and SID fields, respectively, as obtained from the virtual circuit table of the GATEWAY1, such identifications corresponding to the particular virtual circuit and session identified in the VCID and SID fields of the packet 94. The CMD field includes the predefined command GOT_HOST_BY_NAME informing the mobile terminal of the purpose of the packet. The DTA field contains the actual network address(es) of the HOST1 as requested. The packet 96 is transmitted to the mobile terminal 36 which receives the packet and stores the network address.

Step 60 of FIG. 2 relates to establishing the actual socket connection between the mobile terminal 36 and the GATE-

WAY1. Such procedure begins with the mobile terminal 36 generating a packet 98 to be transmitted to the GATEWAY1 as shown in FIG. 4g. The packet 98 includes a header portion directing the packet to the GATEWAY1. In addition, the packet 98 includes the virtual circuit and session identifiers GW1 VCID and GW1 SID for the GATEWAY1 in the corresponding fields. Again, the mobile terminal 36 is able to obtain such information from the corresponding entries in its virtual circuit table (FIG. 5a). With respect to the SOID field, the mobile terminal 36 now looks to its virtual circuit table and selects an available socket corresponding to the previously selected virtual circuit and session. The mobile terminal 36 selects such a socket (nominally labeled MT SOID) and includes it in the SOID field of the packet 98. The CMD field includes a predefined "START_SOCKET" command informing the GATEWAY1 that it is desired to allocate a socket endpoint.

The packet 98 is transmitted to the GATEWAY1 where it is processed as follows. Specifically, the GATEWAY1 selects an available socket corresponding to the information provided in the VCID and SID fields of the packet 98. The GATEWAY1 then stores in association with the selected socket the MT SOID provided in the SOID field of the packet 98. The GATEWAY1 then generates a response packet 100 to be transmitted to the mobile terminal 36 having the format shown in FIG. 4h. Following the standard header portion, the packet 100 includes the MT VCID and MT SID information in the VCID and SID fields similar to the packet 96. In the SOID field, the GATEWAY1 includes the particular socket (nominally labeled GW1 SOID) which was selected from its virtual circuit table (FIG. 5b) in response to the packet 98. Finally, the CMD field includes a predefined command "SOCKET_STARTED" indicating to the mobile terminal 36 that a socket has been established. The packet 100 is then transmitted to the mobile terminal 36.

The mobile terminal receives the packet 100 and proceeds to store the gateway socket information GW1 SOID in its virtual circuit table in association with the mobile terminal socket identified in the SOID field of the packet 98. Hence, after the exchange of the packets shown in FIGS. 4a-4h, the mobile terminal 36 and the GATEWAY1 have established a unique virtual circuit defined by a VCID, SID and SOID combination. The mobile terminal 36 and the GATEWAY1 each has stored in its virtual circuit table the VCID, SID and SOID combination both from the perspective of the mobile terminal 36 and the GATEWAY1.

Next, the mobile terminal 36 attempts to begin actual communications (FIG. 2, step 67) with the particular device (e.g., the HOST1). Initially, however, the mobile terminal 36 generates a packet 102 as shown in FIG. 4i. The packet 102 is addressed to the GATEWAY1 and serves the purpose of requesting that the GATEWAY1 obtain a conventional network connection with the HOST1. Specifically, the packet 102 has a standard header portion addressed to the GATEWAY1. The VCID, SID and SOID fields include the virtual circuit, session and socket information from the perspective of the GATEWAY1, i.e., GW1 VCID, GW1 SID and GW1 SOID, respectively. The CMD field includes the predefined CONNECT command notifying the GATEWAY1 of the desire to establish a connection between the gateway and the HOST1. The DTA field includes the actual network address of the HOST1 as obtained from packet 96 discussed above. It is noted that by this time the mobile terminal 36 has stored in the device entry of its virtual circuit table corresponding to the particular socket the identity of the particular device (e.g., the HOST1) and the gateway address and port corresponding to the GATEWAY1. This facilitates the mobile


terminal 36 keeping track of such information when establishing virtual circuits with several such gateways 40 and devices in the network.

The packet 102 is transmitted to the GATEWAY1 which receives the packet. In response, the GATEWAY1 utilizes conventional network techniques to prepare and/or establish a connection between the GATEWAY1 and the HOST1 as identified in the DTA field of the packet 102. Thereafter, the GATEWAY1 generates a response packet 104 as shown in FIG. 4j. The packet 104 includes the standard header portion in order to be transmitted from the GATEWAY1 to the mobile terminal 36. In addition, the packet 104 includes the corresponding VCID, SID and SOID of the mobile terminal 36 in the appropriate fields. Finally, the CMD field includes a predefined command "CONNECT_RESPONSE" intended to notify the mobile terminal 36 that the connection between the GATEWAY1 and the HOST1 is prepared and/or established.

It is noted that by this time the GATEWAY1 has stored in its mobile terminal entry of its virtual circuit table corresponding to the particular socket, the mobile terminal 36 address, port, and GUID. This facilitates the GATEWAY1 keeping track of which particular mobile terminal 36 is being handled by the respective virtual circuits. In addition, the mobile terminal entry of the virtual circuit table has stored therein the network address of the device (e.g., HOST1) with which the mobile terminal 36 is communicating via the particular virtual circuit connection. Finally, as is discussed below in relation to FIG. 41, the GATEWAY1 will also store in the mobile terminal entry of the table the particular port being used by the GATEWAY1 for communications with the HOST1 with respect to the specific socket connection.

The packet 104 is transmitted to the mobile terminal 36 and the mobile terminal 36 is now prepared to communicate with the HOST1 with the GATEWAY1 serving as an intermediary. For instance, FIG. 4k illustrates a packet 106 containing data which the mobile terminal 36 wishes to transmit to the HOST1. Rather than transmitting the data directly to the HOST1, the data stored in packet 106 is directed via the GATEWAY1 using the socket end points previously established for communications between the mobile terminal 36 and the HOST1. Namely, the header portion again includes the address and port of the GATEWAY1 in the DA and DP fields, respectively. The VCID, SID and SOID fields define the complete connection from the perspective of the GATEWAY1 as obtained from the corresponding entries in the virtual circuit table of the mobile terminal 36. The CMD field includes the predefined command "SEND_DATA" informing the GATEWAY1 that the data is to be delivered to the corresponding device (i.e., the HOST1). The DTA field includes the actual data which is to be delivered to the HOST1.

The packet 106 is then transmitted to the GATEWAY1. The GATEWAY1 receives the packet 106 and in response to the SEND_DATA command looks up the virtual circuit entry in its virtual circuit table corresponding to the information included in the VCID, SID and SOID fields of the packet 106. Based on this information the GATEWAY1 obtains the address information of the HOST1 from the corresponding mobile entry in its virtual circuit table. The GATEWAY1 then generates a packet 108 as shown in FIG. 4i to be sent to the HOST1. The SA field and SP field correspond to the address and port of the GATEWAY1. It is noted that the port utilized by the GATEWAY1 to communicate directly with the HOST1 (or other device) is selected by the GATEWAY1 so as to be unique to the corresponding



socket of the virtual circuit. The DA field and DP field correspond to the address and port of the HOST1 as previously obtained and stored in the virtual circuit table. The DTA field contains the data included in the DTA field in the packet 106 from the mobile terminal 36. It is noted that the packet 108 follows a conventional format native to the end point with which the mobile terminal 36 is communicating.

FIG. 4m represents a packet 110 indicative of a communication sent back to the GATEWAY1 by the HOST1 in response to the packet 108. The DTA field of the packet 110 includes appropriate response data which ultimately is intended for the mobile terminal 36. Since the port identified in the DP field is selected to be unique to a given socket in the virtual circuit table of the GATEWAY1, the GATEWAY1 identifies such socket based on the information stored in the table as discussed above. From the unique gateway port, the GATEWAY1 is able to identify the VCID, SID and SOID of the mobile terminal 36 from the virtual circuit table. The GATEWAY1 then generates a packet 112 as shown in FIG. 4n which is to be transmitted back to the mobile terminal 36. The packet 112 includes the address and port of the mobile terminal 36 in the DA and DP fields, respectively, as obtained from the mobile entry of the virtual circuit table of the GATEWAY1. The DTA field includes the data from the DTA field received in the packet 110 from the HOST1.

Accordingly, communications between the mobile terminal 36 and the HOST1 can remain ongoing via an exchange of packets as represented in FIGS. 4k-4n. The GATEWAY1 simply determines where to direct a received packet based on the contents of its virtual circuit table.

Of course, there is the possibility that the mobile terminal 36 will roam to another LAN (e.g., LAN2) as outlined above. However, as will be apparent such roaming does not have an adverse affect on the virtual circuits formed via the GATEWAY1. Namely, referring back to FIG. 2 assume that the mobile terminal 36 does roam from LAN1 to LAN2. Such roaming is detected in step 72 and the mobile terminal 36 registers with a new access point (e.g., AP2) and obtains a new network identification (e.g., MT ADDRESS* and MT PORT*) (steps 74 and 76). However, the contents of the virtual circuit table (FIG. 5a) of the mobile terminal 36 remain unchanged.

Hence, suppose the mobile terminal 36 wants to continue communicating with the HOST1 in LAN1. The mobile terminal 36 simply generates a packet 114 as shown in FIG. 4o. The source address SA and source port SP fields in the packet 114 will reflect the new network address of the mobile terminal 36 as a result of roaming to LAN2. However, the mobile terminal 36 continues to communicate with the HOST1 via the GATEWAY1 and the specific virtual circuit. Namely, the packet 114 includes the address and port of the GATEWAY1 in the DA and DP fields as shown. In addition, the packet 114, similar to the packet 106, includes the GW1 VCID, GW1 SID and GW1 SOID in the respective fields identifying the particular virtual circuit. The CMD field contains the same "SEND_DATA" command as the packet 106, and the data which is to be transmitted to the HOST1 is included in the DTA field. In addition, though, the packet 114 includes the flag field FLG in which the flag is set to indicate to the GATEWAY1 that the address of the mobile terminal 36 has changed. The packet 114 is then sent to the GATEWAY1 via the access point AP2 and known routing techniques across the WAN 22.

The GATEWAY1 in turn receives the packet 114 and processes the packet in the same manner described above in relation to packet 106 (FIG. 4k) with the following excep-

tion. The GATEWAY1 detects from the FLG field that the address of the mobile terminal 36 has changed. In response, the GATEWAY1 updates its virtual circuit table (FIG. 5b) so as to now include the updated address of the mobile terminal 36 in the corresponding mobile entry. Such updated address is obtained via the SA field and SP field of the packet 114. Otherwise, the GATEWAY1 continues to act as an intermediary and forwards the data in the DTA field to the HOST1 in a packet in the exact same manner as described above in relation to packet 108 (FIG. 4i). In other words, the same procedures described above in relation to FIGS. 4k-4n are repeated. In this manner, the HOST1 is able to continue communicating with the mobile terminal 36, and vice versa, regardless of the fact that the network address of the mobile terminal 36 has changed. In situations where the mobile terminal 36 roams to another LAN but does not need to communicate with the HOST1 immediately, the mobile terminal 36 in the present embodiment is configured to send a gratuitous update packet to the GATEWAY1 in which the FLG field is set so that the GATEWAY1 can immediately update its tables and continue forwarding packets to the mobile terminal 36. Accordingly, seamless roaming is achieved.

FIG. 6 is a block diagram representing the basic structure of the mobile terminals 36 according to the exemplary embodiment. Each mobile terminal 36 includes a processor 170 which can be programmed to control and to operate the various components within the mobile terminal 36 in order to carry out the various functions described herein. The processor 170 is coupled to an operator input device 172 which allows an operator to input data to be communicated to the corresponding LAN such as inventory data, patient information, etc. This information may be sent to the host computer 42 which serves as a central data location, for example, or to a cash register connected to the network backbone 26, as another example, for providing price information. The input device 172 can include such items as a keypad, touch sensitive display, etc. The mobile terminal 36 also may include a bar code scanner 173 coupled to the processor 170 for providing another form of data input. A display 174 is also connected to and controlled by the processor 170 via a display driver circuit 175. The display 174 serves as a means for displaying information stored within the mobile terminal 36 and/or received over the network backbone 26 via an access point 28. The display 174 can be a flat panel liquid crystal display with alphanumeric capabilities, for example, or any other type of display as will be appreciated.

A memory 176 is included in each mobile terminal 36 for storing program code executed by the processor 170 for carrying out the functions described herein. The actual code for performing such functions could be easily programmed by a person having ordinary skill in the art of computer programming in any of a number of conventional programming languages based on the disclosure herein. Consequently, further detail as to the particular code has been omitted for sake of brevity. The memory 176 also serves as a storage medium for storing the above described virtual circuit table for the mobile terminal 36.

Each mobile terminal 36 also includes its own RF section 178 connected to the processor 170. The RF section 178 includes an RF receiver 182 which receives RF transmissions from an access point 28 and via an antenna 184 and demodulates the signal to obtain the digital information modulated therein. An example of a suitable RF receiver 182 for use in the mobile terminal 106 is the Model 025 Direct Sequence Spread Spectrum Radio Module, which is com-

mercally available from Aironet Wireless Communications, Inc. of Akron, Ohio.

The RF section 178 also includes an RF transmitter 186. In the event the mobile terminal 106 is to transmit information to a LAN in response to an operator input at input device 172, for example, the processor 170 forms within the memory 176 the aforementioned information packets. The information packets are delivered to the RF transmitter 186 which transmits an RF signal with the information packet modulated thereon via the antenna 184 to the access point 28 with which the mobile terminal 36 is registered.

Referring now to FIG. 7, a block diagram representative of each access point 28 is shown. Each access point 28 is connected to the network backbone 26 via a connector 240 such as a DB-9 or RJ-45 connector. The connector 240 is connected to the network backbone 26 at one end and to a network adapter transceiver 252 included in the base station 108 at the other end. The network adapter transceiver 252 is configured according to conventional network adapter transceiver techniques to allow the access point 28 to communicate over the network backbone 26. The network adapter transceiver 252 is also connected to an internal bus 254 included within the access point 28. The access point 28 further includes a processor 256 connected to the bus 254 for controlling and carrying out the operations of the access point 28. The processor 256 may include any of a variety of different microprocessors, such as the Motorola 68360 (25 MHz) or Intel 80386 microprocessors.

The access point 28 also includes a memory 258 connected to the bus 254. The memory 258 stores program code executed by the processor 256 to control the other elements within the access point 28 to carry out the functions described herein. It will be readily apparent to a person having ordinary skill in the art of computer programming how to program the processor 256 and the other elements within the access point 28 to carry out the operations described herein using conventional programming techniques based on the flowcharts and descriptions provided herein. As a result, additional detail as to the specific program code has been omitted. Moreover, the memory 258 functions to store information tables maintained by the processor 256 including information such as a list of the mobile terminals 36 which are currently registered with the access point 28.

Also connected to the bus 254 is an RF section 260 included in the access point 28. The RF section 260 includes the aforementioned antenna 32 for receiving radio signals from and transmitting radio signals to mobile terminals 36 within the cell area of the access point 28. Information transmitted from a mobile terminal 36 is received via the antenna 32 and is processed by an RF receiver 262 which demodulates and decodes the signal and converts the information to a digital signal having the aforementioned packet format.

Thereafter, the processor 256 stores the packet in the memory 258 until such time as the base station 108 is able to transmit the information packet onto the network backbone 26 via the network adapter transceiver 252 and connector 240.

Information packets which are transmitted to the access point 28 via the network backbone 26 for transmission to a mobile terminal 36 are received by the network transceiver 252. The processor 256 controls an RF transmitter 264 included in the RF section 260, the RF transmitter 264 also being connected to the bus 254. The processor 256 causes the RF transmitter 264 to modulate an RF signal using

spread spectrum techniques, for example, which in turn carries the information packet to the appropriate mobile terminal 36.

FIG. 8 represents a block diagram of a gateway 40 in accordance with the present invention. Each gateway 40 is connected to the network backbone 26 via a connector 340 such as a DB-9 or RJ-45 connector. The connector 340 is connected to the network backbone 26 at one end and to a network adapter transceiver 352 included in the gateway 40 at the other end. The network adapter transceiver 352 is configured according to conventional network adapter transceiver techniques to allow the gateway 40 to communicate over the network backbone 26. The network adapter transceiver 352 is also connected to an internal bus 354 included within the gateway 40. The gateway 40 further includes a processor 356 connected to the bus 354 for controlling and carrying out the operations of the gateway 40.

The gateway 40 also includes a memory 358 connected to the bus 354. The memory 358 stores program code executed by the processor 356 to control the other elements within the gateway 40 to carry out the functions described herein. It will be readily apparent to a person having ordinary skill in the art of computer programming how to program the processor 356 and the other elements within the gateway 40 to carry out the operations described herein using conventional programming techniques based on the flowcharts and descriptions provided herein. As a result, additional detail as to the specific program code has been omitted. Moreover, the memory 358 functions to store the aforementioned virtual circuit table (FIG. 5b) for the gateway 40.

Although the invention has been shown and described with respect to certain preferred embodiments, it is obvious that equivalents and modifications will occur to others skilled in the art upon the reading and understanding of the specification. For example, while the preferred embodiment utilizes three different layers for the virtual circuit (i.e., VCID, SID and SVID), it will be appreciated that some other number could also be used (e.g., 1, 2, 4, etc.).

Further, the present invention has been described with respect to wireless mobile devices utilizing RF to communicate with devices connected to the network backbones. However, the scope of the present invention also includes mobile devices which do not use RF to communicate but rather are physically connected and disconnected from each respective network backbone via a standard serial or parallel port or other wired network connection, for example.

Additionally, although the present invention was described with a network having discrete components such as the gateway 40, host 42, and DNS 44, it will be appreciated that all of these components could have been combined to form a single unit which can carry on the functionalities described herein.

The present invention includes all such equivalents and modifications, and is limited only by the scope of the following claims.

What is claimed is:

1. A communication system, comprising:

a plurality of local area networks (LANs), each of the LANs including:

a network backbone; and

at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with a device coupled to the network backbone via the at least one access point; wherein when the mobile terminal is registered to at least one access point in one of the plurality of LANs

21

the mobile terminal is assigned a first network address, and when the mobile terminal is registered to at least one access point in another of the plurality of LANs the mobile terminal is assigned a second network address in place of the first network address, the second network address being different from the first network address;

a system backbone interconnecting the plurality of LANs for permitting communications between the plurality of LANs;

a gateway controller, operatively coupled to one of the plurality of LANs, for serving as an intermediary for communications between the mobile terminal and a device on one of the system backbones in order that in the event the mobile terminal is assigned a different network address by virtue of registering with an access point in another of the LANs, the device is able to maintain communications with the mobile terminal without requiring knowledge of a change in the network address of the mobile terminal; and

the gateway controller storing information relating to the current network address of the mobile terminal and the current network address of the device, wherein the gateway controller receives communications from the mobile terminal containing information intended for the device and forwards the information to the device, and the gateway controller receives communications from the device containing information intended for the mobile terminal and forwards the information to the mobile terminal.

2. The communication system of claim 1, wherein each of the plurality of LANs includes a gateway controller connected to the system backbone.

3. The communication system of claim 1, wherein the system backbone represents part of a wide area network (WAN).

4. The communication system of claim 1, wherein the mobile terminal comprises at least one of a data terminal, a telephone, and a pager.

5. A gateway controller for use in a communication system including a plurality of local area networks (LANs), each of the LANs including a network backbone; and at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with a device on the network backbone via the at least one access point, the gateway controller comprising:

means for operatively coupling the gateway controller to the network backbone of a selected one of the plurality of LANs;

a gateway address circuit for communicating with a mobile terminal registered with the access point of the selected LAN via the network backbone for establishing indicia distinguishing the mobile terminal from among other mobile terminals;

a communication circuit for receiving communications from the registered mobile terminal intended for the device and forwarding the received communications to the device together with at least part of the indicia as provided by the communication circuit, and for receiving communications, including the at least part of the indicia, from the device intended for the registered mobile terminal and forwarding the received communications to the registered mobile terminal based on the at least part of the indicia.

6. The gateway controller of claim 5, wherein the indicia represents a layer of addressing utilized in communications

22

involving the gateway controller and the registered mobile terminal but which is not required in communications between other devices communicating on the network backbone.

7. The gateway controller of claim 6, wherein the layer of addressing permits the device wirelessly communicating with the mobile terminal to continue communicating with the registered mobile terminal even if another layer of addressing for the registered mobile terminal changes unbeknownst to the device.

8. The gateway controller of claim 5, wherein the indicia represents a plurality of layers of addressing utilized in communications involving the gateway controller and the registered mobile terminal but which are not required in communications between other devices communicating on the network backbone.

9. The gateway controller of claim 5, wherein the communication circuit includes a memory table in which the indicia for a plurality of mobile terminals are stored.

10. A network communication system, comprising:

a plurality of local area networks (LANs), each of the LANs including:

a network backbone; and

at least one access point coupled to the network backbone which, when a mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with the network backbone via the at least one access point;

a system backbone interconnecting the plurality of LANs for permitting communication between the plurality of LANs;

a gateway controller, operatively coupled to one of the plurality of LANs and able to communicate with devices on other of the plurality of LANs via the system backbone, the gateway controller serving as an intermediary for communications between the mobile terminal and a device on a first of the plurality of LANs; the gateway controller assigning a gateway address to the mobile terminal such that once the mobile terminal starts a session with the device coupled to the first of the plurality of network backbones, the mobile terminal may continuously maintain the session through the gateway even if the mobile terminal registers with an access point coupled to another of the plurality of network backbones; and

the gateway controller storing information in association with the gateway address relating to the current network address of the mobile terminal and the current network address of the device, wherein the gateway controller receives communications from the mobile terminal containing information intended for the device and forwards the information to the device, and the gateway controller receives communications from the device containing information intended for the mobile terminal and forwards the information to the mobile terminal.

11. The network communication system of claim 10, wherein each of the plurality of LANs includes a gateway controller connected to the system backbone.

12. The network communication system of claim 10, wherein the system backbone represents part of a wide area network (WAN).

13. The network communication system of claim 10, wherein the mobile terminal comprises at least one of a data terminal, a telephone, and a pager.

14. A mobile terminal for use in a communication system including a plurality of local area networks (LANs), each of

23

the LANs including a network backbone and at least one access point coupled to the network backbone which, when the mobile terminal is registered to the access point, enables the mobile terminal to communicate wirelessly with a device coupled to the network backbone via the at least one access point, the mobile terminal comprising:

mobile terminal address circuitry for establishing with a gateway controller coupled to one of the plurality of LANs via the network backbone indicia distinguishing the device from among other devices; and

a communication circuit for transmitting communications destined for the device to the gateway controller together with at least part of the indicia, and for receiving communications transmitted by the device and destined for the mobile terminal from the gateway controller, the received communications having the at least part of the indicia included by the gateway controller.

15. The mobile terminal of claim 14, wherein the indicia represents a layer of addressing utilized in communications between the mobile terminal and gateway controller for identifying the device.

16. The mobile terminal of claim 15, wherein the layer of addressing permits the device communicating with the mobile terminal to continue communicating with the mobile terminal through the gateway even if another layer of addressing of the mobile terminal changes unbeknownst to the device.

17. The mobile terminal of claim 14, wherein the mobile terminal address circuitry includes a memory table in which the indicia for a plurality of devices are stored.

24

18. A mobile terminal for use in a wireless communication system including a plurality of local area networks (LANs), each of the LANs including a network backbone and at least one access point coupled to the network backbone, the mobile terminal including:

a portable housing;

means for assigning unique indicia to a plurality of devices in the wireless communication system with which the mobile terminal desires to communicate; and means for storing the unique indicia for the plurality of devices.

19. The mobile terminal of claim 18, further comprising a transceiver for wirelessly transmitting and receiving communications with at least one of the plurality of devices via a gateway controller coupled to the network backbone.

20. The mobile terminal of claim 19, wherein the communications received by the gateway controller from the one of the plurality of devices includes at least part of the unique indicia assigned to the one of the plurality of devices by the mobile terminal.

21. The mobile terminal of claim 20, wherein the unique indicia represents a layer of addressing utilized in communications between the mobile terminal and gateway controller.

22. The mobile terminal of claim 18, wherein the means for storing is a memory.

* * * * *